

Immersive Search: Using Virtual Reality to Examine How a Third Dimension Impacts the Searching Process

Austin R. Ward

School of Information and Library Science
University of North Carolina - Chapel Hill
austinrw@unc.edu

Rob Capra

School of Information and Library Science
University of North Carolina - Chapel Hill
rcapra@unc.edu

ABSTRACT

In this paper, we present results from an exploratory study to investigate users' behaviors and preferences for three different styles of search results presentation in a virtual reality (VR) head-mounted display (HMD). Prior work in 2D displays has suggested possible benefits of presenting information in ways that exploit users' spatial cognition abilities. We designed a VR system that displays search results in three different spatial arrangements: a list of 8 results, a 4x5 grid, and a 2x10 arc. These spatial display conditions were designed to differ in terms of the number of results displayed per page (8 vs 20) and the amount of head movement required to scan the results (list < grid < arc). Thirty-six participants completed 6 search trials in each display condition (18 total). For each trial, the participant was presented with a display of search results and asked to find a given target result or to indicate that the target was not present. We collected data about users' behaviors with and perceptions about the three display conditions using interaction data, questionnaires, and interviews. We explore the effects of display condition and target presence on behavioral measures (e.g., completion time, head movement, paging events, accuracy) and on users' perceptions (e.g., workload, ease of use, comfort, confidence, difficulty, and lostness). Our results suggest that there was no difference in accuracy among the display conditions, but that users completed tasks more quickly using the arc. However, users also expressed lower preferences for the arc, instead preferring the list and grid displays. Our findings extend prior research on visual search into the area of 3-dimensional result displays for interactive information retrieval in VR HMD environments.

KEYWORDS

immersive search, virtual reality, three-dimensional search

ACM Reference Format:

Austin R. Ward and Rob Capra. 2020. Immersive Search: Using Virtual Reality to Examine How a Third Dimension Impacts the Searching Process. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*, July 25–30, 2020, Virtual Event, China. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3397271.3401303>

SIGIR '20, July 25–30, 2020, Virtual Event, China

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in: ACM ISBN 978-1-4503-8016-4/20/07...\$15.00
<https://doi.org/10.1145/3397271.3401303>

1 INTRODUCTION

Interactive information retrieval (IIR) researchers have studied the effects of search engine result presentation in two-dimensional interfaces for years. However, less research has examined results presentation in the types of 3D spaces enabled by emerging virtual reality (VR) head-mounted displays (HMDs). Prior work has explored interfaces that arrange information in 3 dimensions to assist in navigating information spaces [8] and that imply relevancy through distance [4, 9]. However, these prior studies have been limited by the use of traditional 2D desktop monitors that can only use visual cues to engage users' spatial cognition (the mechanism through which humans acquire, process, and utilize information about a space). Immersive technologies, such as virtual reality head-mounted displays (VR HMDs), can increase spatial cognition by activating users' vestibular and proprioceptive systems [1]. Research with information displays in VR has shown that spatial arrangements of information can aid with recall and simple visual search [2, 6]. However, few efforts have been made to explore how spatial cognition engaged through immersive technologies could benefit IIR.

We report on a laboratory study (N=36) to explore the effects of different spatial arrangements of search results in an immersive virtual environment for visual search tasks. Participants completed 18 visual search trials across three different spatial display conditions within a custom-built application for the Oculus Quest VR headset. Each trial asked the user to either locate a relevant search result from a set of 40 results, or to determine that no relevant result was present in the set. Each spatial display condition used a different arrangement of results: (1) a 1x8 list, similar to what modern search engines provide; (2) a 4x5 grid,

that required some head movement; and (3) a 2x10 arc, a wide arrangement where twisting the body and head were required to view all results. The study investigated these research questions:

RQ0: What is the effect of target presence on participants' search? This was intended as a manipulation check to verify that the absence of a relevant target led to longer search times.

RQ1: What are the effects of spatial display and target presence on participants' search interactions?

RQ2: What are the effects of spatial display and target presence on participants' perceptions about the search task and interface?

2 RELATED WORK

Prior work in the area of IIR and spatially arranged information has suggested that users can understand and utilize spatial arrangements of data [4, 8, 9]. These studies used 2D, non-immersive displays, but showed that participants were able to perform retrieval

tasks more quickly and with fewer incorrect results in certain conditions. Immersive virtual environments (IVEs) could further leverage humans' spatial cognitive abilities through engaging their vestibular and proprioceptive systems [6].

Information search in immersive HMDs has been studied in the context of simple visual search tasks. Pausch, Proffitt, and Williams [7] found that participants in a virtual room were significantly faster at determining that a search target was not present when using an immersive display versus a traditional 2D display. Billinghurst, Bowskill, Dyer, and Morphett [2] conducted an experiment in which participants looked for targets in three display conditions: a) a non-spatial view using a trackball to change the current page of results, b) a spatial view in which the results conceptually surrounded the user and the currently focused page was changed using a trackball, and c) a view similar to (b) except that page changes were accomplished via head movement, with a single page visible at a time across all conditions. Participants performed significantly faster on visual search tasks in the spatial conditions (b and c). In addition, participants expressed higher satisfaction for the spatial conditions and rated them the highest in terms of understanding where the information was.

Other research has shown the benefits of IVEs and spatial cognition in memory [3, 5, 6]. Krokos, Plaisiant, and Varshney [6] found that tasks designed to engage vestibular and proprioceptive systems enhanced participants' recall and reduced errors versus traditional displays. In a study about navigation [1], participants that used immersive technology showed more accurate knowledge of a virtual environment than those that did not.

3 METHODS

We conducted a user study with 36 participants (22 female) to investigate our research questions. Participants were recruited from students and employees at our university through an opt-in mailing list. A majority of participants (26) reported having spent ≤ 30 minutes ever using VR systems before.

Display conditions – We used a within-subjects, repeated measures experimental design to compare three display conditions: 1) list – a vertical list of 8 search results aligned in a flat plane in front of the participant; 2) grid – a 4x5 array of search results curved around the participant's forward view (140 degrees); and 3) arc – 2 rows of 10 results that surrounded the participant across a 220 degree arc (see Figure 1). These conditions were selected to compare a condition similar to current 2D SERPs (list) against a condition that required head movement (grid) and a condition that required both head and body movement (arc).

Result items were presented as text, formatted similarly to traditional 2D search result surrogates, including a title, URL, and snippet (Figure 2). Each display condition presented a total of 40 search results. To keep the search surrogates of similar size and meet the movement requirements, the list displayed 8 results per page (5 total pages) while the grid and arc each displayed 20 results per page (2 total pages). All conditions were presented in a VR HMD (Oculus Quest).

Tasks – Participants were asked to complete one practice and six search trials in each of the three display conditions. In each trial, participants were given a target description and a set of results




Figure 1: A top-down view of the display conditions.



Figure 2: Search surrogate example.

presented according to the display condition (list, grid, arc). Participants were instructed to find the target result or to indicate that the target was not present. For example, one of our trials asked, "From the given search results, find the result that will help you answer 'What color is a giraffe's tongue.' Find and select the relevant search result or indicate that no relevant result is available." Participants did not issue queries and could not view the landing pages for the results, so they had to make their decisions based only on the information shown in the result set. Participants could select a specific result by moving their head to focus a highlight on the desired result and then confirm their selection by pressing a button on a hand-held controller. The Oculus Quest system we used does not support eye tracking, but instead can interpret head movements to control which result is highlighted. Thus, it was possible (to a limited degree) for a participant to keep the focus on one item (fixed head position) while looking at nearby items (changing eye gaze position). Pagination was supported by two buttons on the controller - a page forward button and a page back button. Participants were seated in a stationary chair for all trials.

Half of the trials for each task topic were populated with a single, clearly relevant target and the other half contained no relevant results. The order of trials, the display conditions, and the presence of a target relevant result were counter-balanced using balanced Latin squares. Relevant and non-relevant results were drawn from Bing, using related or unrelated topic queries, respectively. Similar visual search tasks have been used in prior work on spatial displays [2, 7]. We choose this type of visual search task for several reasons: 1) determining if a relevant result is present is an important component of post-query SERP results examination, 2) we wanted to avoid possible confounding factors (e.g., users' judgements about which of several possible results might be the "most relevant"), and 3) so that we could compare our results with prior work. Future work should examine more realistic scenarios with multiple relevant results.

System – The system application was built for the Oculus Quest VR HMD platform using the Unity game engine and Unity primitives. The application logged all system interactions to an external database and used Oculus APIs for tracking head movements and head direction.




Figure 3: User Perception Measures

4 MEASURES

Interaction Data: The system recorded interaction data for focus events and click events. *Focus events* were logged when the participant moved their head to highlight a search result. As described previously, the Oculus Quest does not include eye tracking, so the focus events we logged were based on head movement. *Click events* were logged when a participant pressed the controller trigger to select a highlighted item. For each trial, we recorded the time to complete the trial, the answer selected, the correct response, and attributes of the target (presence, position, and page number).

User's perception data: During the study, participants were asked to complete several questionnaires about their perceptions of the system and tasks. After completing each set of 6 trials in a given display condition, participants were given a *post-condition questionnaire* with two sections designed to collect participants' immediate reactions to using each of the display conditions. The first section included 4 questions (7-point scales from 1=very low to 7=very high) about the amount of mental demand, physical demand, workload, and difficulty determining when a correct result was not available (1=very easy, 7=very hard). The second section included a set of 7 statements with Likert-type agreement scales (1=strongly disagree, 7=strongly agree): 1) Overall, I felt this task was easy; 2) I felt generally uncomfortable while using this display condition; 3) I felt eyestrain...; 4) I felt dizzy...; 5) I am satisfied with the amount of time it took to complete this task; 6) I thought the arrangement of the search results made this task easy to complete; 7) I felt confident in the answers I selected.

After completing all 18 trials, participants were given a final *post-session questionnaire*. The post-session questionnaire was designed to gather participants' perceptions *after experiencing all three display conditions*. It included a set of 7 statements with Likert-type agreement scales (1=strongly disagree, 7=strongly agree): 1) It was difficult to find the correct result; 2) It was difficult to determine when a correct result was not present; 3) It was hard to use this display condition; 4) I felt this display condition made the task take too long; 5) This display condition required a lot of backtracking; 6) This display condition made me feel lost; 7) This display condition made it hard to feel confident in my answers. The questionnaire also included open-ended questions asking participants to list things that they liked and disliked about each display condition by typing responses into a text box, and to rate their overall satisfaction with each of the three display conditions on a scale from 0 (low) to 100 (high).

5 RESULTS

5.1 RQ0: Target Presence Manipulation Check

RQ0 was intended as a manipulation check to make sure that target presence impacted focus events and trial completion times. As expected, ANOVAs showed that target presence had a significant effect on the number of focus events ($F(1, 642) = 278.66, p < .000$) and trial completion time ($F(1, 642) = 271.61, p < .000$). When the target was not present (M_0) compared to when it was present (M_1), there were on average more focus events ($M_0=54.82$ vs. $M_1=26.26$) and longer completion times ($M_0=51.64$ vs. $M_1=26.19$). These results are consistent with our expectation and results from Pausch et al.[7].

5.2 RQ1: Effects on Interaction Data

RQ1 considers the effect of display condition and target presence on users' search behaviors. Across all 648 trials (36 participants x 18 trials), only 12 were answered incorrectly (all 12 were false-negatives; 5 in the list condition, 5 in the grid, and 2 in the arc). A chi-squared test did not find any significant difference among these ($\chi^2(2, N = 648) = 1.5283, p = .47$). Since there were only 12 incorrect trials, we excluded them from the following analysis.

Two-way ANOVA did not find significant main or interaction effects of display condition and target presence on focus events. However, we did find an interaction effect on trial completion time ($F(2, 642) = 3.30, p < .037$). When the target was present, no significant differences in trial completion time were found across the display conditions. However, when the target was not present, there were significant differences in trial completion times ($F(2, 324) = 6.19, p < .01$). Post-hoc pairwise tests (with Tukey correction) showed that when the target was not present, participants completed trials more quickly in the arc ($M=47.06, SD=19.55$) and grid ($M=50.22, SD=24.68$) display conditions compared to the list ($M=57.64, SD=23.55$).

5.3 RQ2: Effects on Users' Perceptions

RQ2 considers the effect of the display conditions on participants' perceptions about the search task and interface. Recall that users' perceptions were measured using the post-condition and post-session questionnaires, so our analysis examines the main effect of display condition. One-way repeated-measures ANOVAs (with Tukey post-hoc comparisons) found significant main effects of display condition on participants' perceptions as shown in Figure 3. Below, we summarize the RQ2 results by considering comparisons between the list, grid, and arc conditions.

List vs Arc – The list condition was the closest to a traditional 2D SERP, and was reported to be significantly preferred to the arc condition across several dimensions: 1) physical demand, 2) uncomfortableness, 3) the ease of the arrangement of the results, 4) difficulty determining when a relevant result was present, 5) being hard to use, and 6) making it hard to feel confident in my answers.

Grid vs Arc – The grid condition used more of the space immediately around the user and only required head movements to view. It was reported to be significantly preferred to the arc condition across several dimensions: 1) physical demand; 2) eyestrain, 3) ease of the arrangement of the results, 4) difficulty determining when a relevant result was present, 5) being hard to use, 6) making the task take too long, 7) requiring a lot of backtracking, 8) feeling lost, and 9) making it hard to feel confident in my answers.

List vs Grid – Very few significant differences in perceptions between the list and grid conditions were found. Only physical demand showed a significant difference.

Overall preference – The final question in the post-session questionnaire asked participants to score each display condition on a scale of 1-100. The participants significantly scored the list and grid higher than the arc ($M_{list}=71.92$, $M_{grid}=80.86$, $M_{arc}=48.97$).

5.4 Written Responses

In the post-session questionnaire, we asked participants to list things that they liked and disliked about each display condition. The primary author qualitatively coded their responses using two-rounds of inductive coding. In their responses, participants noted liking the list for requiring less head movement and for its familiarity. However, they mentioned disliking the list for requiring more pagination and for the amount of backtracking required to feel confident that they have completed a thorough search. Participants noted liking the grid and arc conditions for the higher number of results on each page (20), that they required less pagination, that there was more space between results (mentioned as making browsing easier), and for being more interesting than the 2D list. Some participants mentioned disliking the grid for being overwhelming and having no clear scan path. Others wished the grid had contained one less row or column of results. Interestingly, the grid was the only condition for which several participants explicitly indicated no dislikes. For the arc, participants noted disliking the amount of head and body movement required, especially for results on the extreme right and left of the arc.

6 DISCUSSION

Our analysis shows several interesting results. First, our RQ1 analysis did not find significant differences across the display conditions when the target was present. However, when the target was not present, participants completed trials more quickly in the arc and grid compared to the list. Second, our RQ2 results show that across a variety of measures, participants perceived the arc condition more negatively than the list and grid. Third, our qualitative results indicated that participants were familiar with the list layout, liked the number of results displayed in the grid, and felt that the arc required a lot of physical movement. Finally, we note that across all three display conditions, there were very few incorrect results, illustrating that participants were able to complete the tasks accurately.

Our results are in line with prior work. In Pausch et al. [7], participants were placed in a virtual room filled with letters arranged on the walls in two display conditions (Head-tracked HMD vs stationary display) and tasked with either finding a target letter or determining that the target letter was not present. In the searches where the target was present, there was no significant difference between conditions in the time to find the target. However, participants completed their searches significantly faster when the target was absent in the head-tracked HMD condition versus the stationary display. This aligns with our findings from the interaction data which showed participants performing target absent trials quicker in the grid and arc spatial display conditions than in the list.

Participant perceptions on the use of each condition showed an affinity for the list and grid conditions over the arc. However, the average completion time for target absent trials was significantly lower in the grid and arc conditions than the list. The qualitative responses provide possible reasons for this divergence. Participants noted that the arc was too wide and that they disliked the body twisting required to view all of the results in the arc. However, these types of comments were not noted as often for the grid. This suggests that users may have a threshold of movement they are willing to perform for visual search tasks.

How to most effectively arrange search results for interactive information retrieval in immersive virtual environments requires further study. Our findings suggest that: 1) users are capable of utilizing non-traditional arrangements of information that leverage the third dimension, 2) the grid and arc arrangements led to faster “complete” scans of the information space (e.g., the target absent condition), and 3) among the spatial displays, users expressed preferences for the grid over the arc. Our results suggest that immersive search interfaces should consider a possible trade-off between familiarity and a time cost associated with pagination. However, our results with the arc condition suggest that displaying results in too wide a field (e.g., 220 degrees) may require more body movement than users prefer. Future work should consider more complex tasks that involve multiple relevant results in the display.

Acknowledgements. This work was supported in part by NSF grants #1718295 and #1552587. Any opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the NSF.

REFERENCES

- [1] Niels H. Bakker, Peter O. Passenier, Peter J. Werkhoven, Henk G. Stassen, and Peter A. Wieringa. 2001. Navigation in Virtual Environments. *IFAC Proc. Volumes* 34, 16 (2001), 89 – 94.
- [2] Mark Billinghurst, Jerry Bowskill, Nick Dyer, and Jason Morphet. 1998. An evaluation of wearable information spaces. In *Proc. IEEE 1998 Virtual Reality Annual International Symposium*. IEEE, 20–27.
- [3] Joel Harman, Ross Brown, and Daniel Johnson. 2017. Improved memory elicitation in virtual reality: new experimental results and insights. In *IFIP Conference on Human-Computer Interaction*. Springer, 128–146.
- [4] Matthias Hemmje, Clemens Kunkel, and Alexander Willett. 1994. LyberWorld—a visualization user interface supporting fulltext retrieval. In *SIGIR’94*. 249–259.
- [5] Jan-Paul Huttner, Kathrin Robbert, and S. Robra-Bissantz. 2019. Immersive Ars Memoria: Evaluating the Usefulness of a Virtual Memory Palace. In *Proc. HICSS*.
- [6] Eric Krokos, Catherine Plaisant, and Amitabh Varshney. 2019. Virtual memory palaces: immersion aids recall. *Virtual Reality* 23, 1 (2019), 1–15.
- [7] Randy Pausch, Dennis Proffitt, and G. Williams. 1997. Quantifying immersion in virtual reality. In *Proc. Conf. Comp. Graphics and Interactive Techniques*. 13–18.
- [8] George Robertson, Mary Czerwinski, and Maarten Van Dantzich. 1997. Immersion in desktop virtual reality. In *Proc. UIST*. 11–19.
- [9] Tuukka Ruotsalo, Jaakko Peltonen, Manuel JA Eugster, Dorota Glowacka, Patrik Florén, Petri Myllymäki, Giulio Jacucci, and Samuel Kaski. 2018. Interactive intent modeling for exploratory search. *ACM TOIS* 36, 4 (2018), 1–46.