Virtual Memory Palaces: Immersion aids Recall

Eric Krokos*and Amitabh Varshney†

Department of Computer Science and Institute for Advanced Computer Studies
University of Maryland College Park

Abstract

Virtual reality displays, such as head-mounted displays (HMD) afford us a superior spatial awareness by leveraging our vestibular and proprioceptive senses. Since the classical times, people have used memory palaces as a spatial mnemonic to help remember information by organizing it spatially in an environment and associating it with salient features in that environment. In this paper, we explore whether using virtual memory palaces in a head mounted display (HMD) will allow a user to be able to better recall information than when using a traditional desktop display. Our study shows that virtual memory palaces coupled with increased immersion in a HMD provide us with a superior memory recall compared to traditional desktop displays. This, we hope, is a first step in using virtual environments for creating more memorable experiences that enhance productivity through better recall of large amounts of information organized around the idea of virtual memory palaces.

Keywords: Immersion, experimental methods, HMD, 3D navigation, visualization, psychology, training, education, user-study, perception, presence

1 Introduction

Since the classical times, people have used memory palaces (method of loci), by taking advantage of the brain's ability to spatially organize thoughts and concepts. In a memory palace, one mentally navigates an imagined structure to recall information [4]. Even the Roman orator Cicero is believed to have used the memory palace technique by visualizing his speeches and poems as spatial locations within the auditorium he was in. Virtual reality displays, in contrast to traditional displays, can combine immersive spatial representations of data with natural interactions. Our ability to generate data has far surpassed our ability to comprehend it [10]. Virtual reality offers us a way to organize large information spaces in the form of virtual memory palaces that will leverage our natural abilities to understand, navigate, and recall in our quest to better understand the data and its relationships. In this paper we present the results of a user study that involved building virtual memory palaces, and testing to see if they could assist in superior recall aided by the immersion afforded by a head-mounted display (HMD) as compared to using a traditional desktop display. Our results indicate a statistically significant improvement in recall when using virtual memory palaces in immersive, head-tracked HMDs compared to traditional desktops.

*ekrokos@umiacs.umd.edu
†varshney@umiacs.umd.edu

2 Related Work

Memory palaces have been used since the classical times to aid recall by using spatial mappings and environmental attributes. Figure 1 shows a depiction of a memory palace attributed to Giulio Camillo in 1511. The idea was to map words or phrases onto a mental model of an environment (in this case an amphitheater), and then recall those phrases by mentally visualizing that part of the environment.

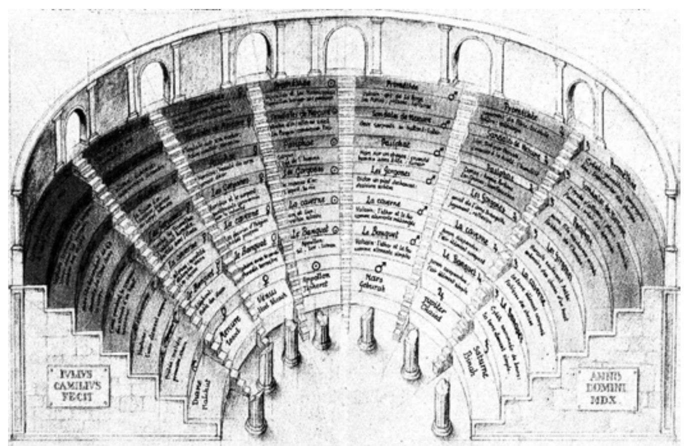

Figure 1: Giulio Camillo's depiction of a memory palace (1511 AD). Memory palaces like this have been used since the classical times as a spatial mnemonic.

2.1 Memory Palaces

Legge *et al.* [6] has compared traditional mental memory palaces (method of loci) against virtual memory palaces. In his study, the subjects were divided into three groups. The first group was instructed to use a mental memory palace, the second group a virtual memory palace, and the third (control) group was not informed on the use of any mnemonic device. The group that was asked to use a virtual memory palace was given five minutes to explore the virtual palace using a mouse and a keyboard. The subjects in the three groups were given 10 to 11 uncorrelated words and asked to memorize the words with their mnemonic device, if any. The users then had to recall the words serially. This user study found that the users who used a virtual memory palace performed better than those using a mental memory palace, and those who were not instructed on a memory strategy did not perform as well as those who were. These results were statistically significant. We note that in the virtual memory palaces used in this study, the participants were expected to imagine the words associated with parts of the virtual memory palace, and that they did not actually *see* the words rendered in that environment. Also, the subjects navigated the virtual memory palace on a desktop and did not experience it in an immersive virtual reality environment.

2.2 Spatial Organization

Spatial organization in a 3D virtual environment has been used to assist in recall. Robertson *et al.* [14] created an application called the *Data Mountain* that allowed users to place documents on a 2D inclined plane in 3D space, using a traditional mouse interface. This inclined plane was rendered using a perspective projection. They found that the users were able to spatially recall the locations and store documents with reduced times, as well as minimize retrieval failure, when using the 2D inclined plane in 3D, instead of the conventional approach of laying them out on a flat 2D display surface. Although this study used 3D rendering, it restricted the viewer to a single view. Interestingly, other studies such as by Cockburn [3] show that there is no statistically significant difference in spatial memory tasks in 2D and 3D visualizations. We note that these studies have not been carried out in immersive environments.

2.3 Display Immersion

Johnson *et al.* [5] performed a study evaluating spatial memory based on two setups: the first used a traditional monocular desktop display and the second used a stereoscopic head-tracked desktop display. The participants in the study were placed in a virtual city and guided through a set of predefined waypoints. The waypoints were then removed and the participants were asked to go through a portion of that path without the waypoints. The study found a statistically significant improvement in episodic memory for the head-tracked stereo display compared to the traditional desktop display. Ragan *et al.* [12] carried out a user study in which participants were asked to memorize and recall the sequence of placement of virtual objects on a grid. The participants were divided into multiple groups that were shown the same process with different fields of view. They found that higher fields of view produced a statistically significant performance improvement. Efficacy of varying immersion levels by changing the field of view has also been studied in the context of procedural training [1]. In a study by Sowndararajan *et al.* [15] they compared subject performance for a task, but with two different fields of view – one with a laptop and the other with a large rear-projected L-shaped display. The study had participants trained on two procedures and the performance with the two levels of immersion was compared. The study found that higher levels of immersion (in this case, field of view) were more effective in learning complex procedures that reference spatial locations. The studies by Johnson *et al.*, Ragan *et al.*, and Sowndararajan *et al.* examine the relationship between memory and immersiveness as measured by stereo, field of view, or head tracking, but not inside a head-mounted display (HMD).

Pausch *et al.* [11] studied if immersion in a HMD aids in searching and detection of information. In their study, they created a virtual room with letters distributed on walls, ceiling, and floor. A user was placed in the center of this room and was asked if a set of letters was present or not. The test was conducted using a HMD and a traditional display with a mouse and keyboard. They found that when the search target was present, the HMD and the traditional display had no statistically significant difference in performance. However, when the target was not present, the users were able to confirm that target was missing more quickly in the HMD than the traditional display. In addition, the users that used the HMD first and then moved to a traditional display had better performance than those who used the display first and then the HMD. This suggests a positive transfer effect from HMD to a desktop.

Mania and Randell [8] examined the productiveness amongst a stereoscopic HMD, a monoscopic HMD, a desktop, and the real world. For this study they recreated a small room, as precisely as possible, in 3D graphics. The users were shown a set of shapes(placed in the real room as well as in the virtual room) for a period and the users then had to recall what shapes were placed where. During the recall phase, the amount of time the participants were idle was measured and not the recall accuracy. This was exhibited by no head or mouse movement. The study found that the users in the HMD had significantly higher idle time during the recall than those in the real world and the other display modalities.

In another study by Mania *et al.* [9] they examined the accuracy and confidence level associated with recall in a similar room as described above. The users were exposed to objects within a room either in a HMD, on a desktop, or in the real world. Then immediately after and one week later, the participants recalled the objects and their locations as well as their confidence in their answers by writing out their answers on a piece of paper showing a drawing of the room. The study found that initially and after one week, participants had the most accurate recall in the real-world scene, were slightly less accurate and confident in the HMD, and least accurate and confident on the desktop. However, results showed that the proportion of correct answers for the most confident score was given in a HMD when participants used the mouse to look around rather than those who used head tracking. An interesting result of this study is that natural-to-senses head tracking does not always correspond to higher accuracy or confidence in recall.

Another study by Mania and Chalmers [7] evaluated the memory recall and confidence of participants after a 15-minute seminar delivered on a desktop, a HMD, or in the real world. This study focused on episodic memory rather than just spatial memory. A virtual replica of the seminar was recreated as accurately as possible to the real-world scenario. The user study found that participants were better able to recall information in the real world compared to the HMD and that there was no statistical difference between the real world and desktop in terms of recalling the events of the seminar. However, they found that the probability of a participant giving a high-confidence answer in the HMD was higher than the desktop, which they concluded could mean that the visual experiences in the HMD are more memorable and vivid.

2.4 Embodied Interaction

Wraga *et al.* [16] compared the effectiveness of vestibular and proprioceptive rotations in assisting recall. They found that users in a HMD were able to better recall by rotating their heads as compared to using a joystick for an equivalent scene rotation. They also found that users in a HMD who controlled their bearing in a virtual world by actively rotating in a swivel chair were able to recall better than those that were being rotated by a tester.

Brooks *et al.* [2] studied if active participants had superior recall of a virtual environment on a desktop than passive participants. Active participants controlled camera navigation via a joystick themselves, where passive participants observed the navigation. They found that those who were actively participating more accurately recalled the environment layout compared to those who were passive. However, they also found that there was no statistically significant difference between the recall or recognition of the objects or their positions within the environment between the active and passive participants. This suggests that memory was only enhanced for those aspects of the environment that were interacted with directly - particularly the environment which was navigated.

Riecke *et al.* [13] studied the importance of full-body walking to achieve effective VR navigation in a HMD. They had participants navigate by (a) using a joystick or (b) by walking, or (c) a hybrid navigation where rotation was done by the body and translation by the joystick. They found that (b) walking was superior to (a) joystick navigation. However, they also found that (b) walking compared to (c) hybrid navigation had roughly equal performance, which suggests that the easier-to-implement hybrid navigation strategies are as good as the full walking navigation.

3 Our User Study

Previous work has examined the role of spatial organization, immersion, and interaction in assisting recall. The goal of our user study is to examine if a virtual memory palace, experienced immersively in a head-tracked stereoscopic head-mounted display, can assist in recall better than a mouse-based interaction on a traditional, non-immersive, monoscopic desktop display. We wish to examine if virtual memory palaces experienced immersively in a HMD can assist in more accurate recall than on a traditional desktop.

For this study we used a traditional desktop with a 30" monitor and an Oculus DK2 HMD. The rendering resolution for both the Oculus DK2 and the desktop was set to 1080×1200 with a rendering field of view (FOV) of 100° degrees.

We had 40 participants in our user study, of which 30 were male and 10 were female, each with normal or corrected-to-normal vision. The study session for each participant lasted around 45 minutes. First, a participant was shown the images of all the 42 celebrity faces used in the study, with names, to familiarize them with who they would later be asked to recall. If a participant could not remember the name of a celebrity during the recall phase, it was considered acceptable for them to unambiguously describe the traits of the celebrity. The celebrities were carefully chosen to be easily recognizable and familiar. Next, the participant would be placed either in front of a desktop monitor with a mouse or inside a head-tracked HMD. The participants were given as much time as they desired to get comfortable with the desktop and HMD display. The users rotated the scene on a desktop monitor using a traditional mouse and in the HMD they rotated their head and body.

Once the participants were comfortable with the setup and the controls, they were shown a set of 21 faces distributed around a 3D scene. We used two such scenes – a palace and a medieval town, shown in Figure 2. The participants were given five minutes to memorize the faces and their locations within the scene. After the five minute period, the display went blank and the participants were given a two-minute break. We added this break to avoid recall from the participants' short-term memory.

After two minutes, the scene would reappear on the display with numbers having replaced the faces, as shown in Figure 3. The participants were then asked to recall which celebrity face had been at each numbered location. During this testing phase, the participants were able to look around and explore the scene just as they did in the training phase, using the mouse on the desktop or rotating their head-tracked HMD. For each numbered location in the scene, the participants recalled the celebrity at that location, as well as a confidence rating for their answer, ranging from 1 to 10, with 10 being certain. The participants had up to five minutes to recall all the celebrities in the scene. Once the participant was confident in all their answers, or the five-minute period had passed, the testing phase ended. After a break, the participants were placed in the other display that they had not previously tested with. The process was then repeated with a different scene and a different set of 21 faces to avoid information overlap from the previous test. At the end, each participant was tested on the two displays, a desktop and a HMD, on two different scenes, and with two different sets of 21 faces.

(a)

(b)

Figure 2: Virtual Memory Palaces used in our user study (a) an ornate palace, and (b) a medieval town

4 Results

Our hypothesis is that a virtual memory palace experienced in an immersive head-tracked head-mounted display will lead

Figure 3: Virtual memory palace: recall phase

to superior recall than on a mouse-controlled desktop display. In our study we had the participants alternate between the two sets of faces, the two scenes, and the starting order of the two displays. The reason for swapping the faces and scenes between displays was to avoid any cross contamination. To mitigate any learning behavior from the first trial to the second, the users alternated which display was used first. By alternating amongst the displays, the scenes, and the faces, we expect to mitigate any confounding effects.

4.1 Accuracy

The primary goal of the study was to examine the recall accuracy differences between the two displays. In Figure 4 we present the overall performance of the users in the HMD compared to the desktop.

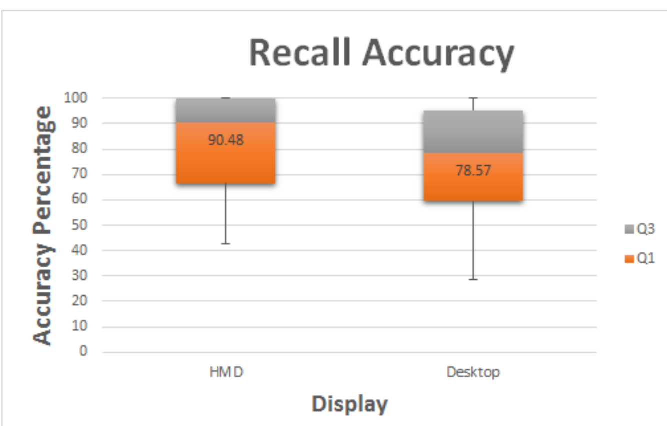

Figure 4: The overall recall performance of participants using a HMD is about 10% higher compared to a desktop. The median recall accuracy percentage for HMD is 0.9048 and for desktop display is 0.7857. The figure shows the 25^{th} to 75^{th} percentiles for each display modality.

This shows that participants in the HMD were better able to recall than those on a traditional desktop. Using a paired t-test, we calculated $p = 0.0017 < 0.05$. Using ANOVA we calculated $F(1, 78) = 4.17$ with $p = 0.0446$. Both show that our result is statistically significant. In addition, we found that the ordering of faces or ordering of scenes had no statistically significant impact on the recall accuracy.

4.2 Confidence

Previous work [7,9] had examined confidence with recall accuracy. This previous work allows us to study not only the objective recall accuracy but also the subjective certainty of the user answers. We had asked the participants to indicate their confidence on a scale of 1 to 10, with 10 being certain, for each answer. The confidence scores aggregated across all the 40 participants and all the 42 faces that each studied are shown in Figure 5.

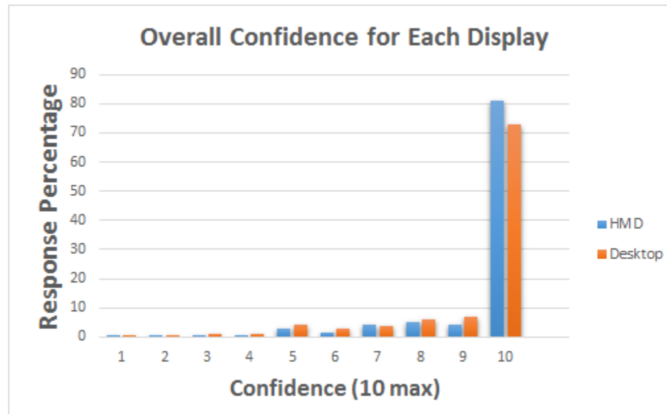

Figure 5: The overall confidence scores of participants using a HMD and a desktop. Each participant gave a confidence score between 1 and 10 for each face they recalled. Those in the HMD are slightly more confident about their answers than those on the desktop.

From the figure above, we can see that users were slightly more confident in the HMD than using a desktop. However, confidence is not always an indication of correctness. We want to see if the HMD was giving a false sense of confidence. Figure (Figure 6) shows the overall correctness of answers given in each display based on the confidence of participant answers.

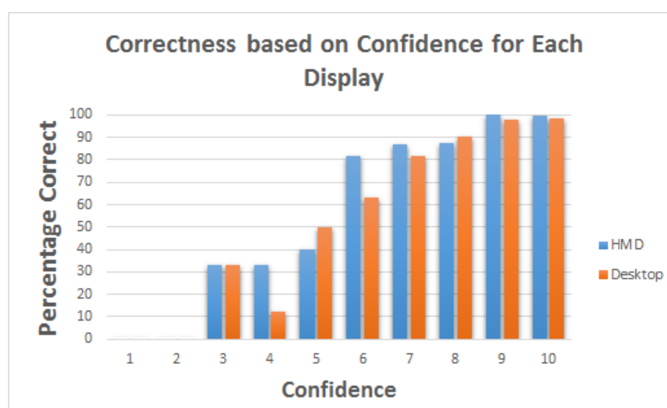

Figure 6: The percentage of times participants gave the correct answer given their confidence score.

The results above show that when the users were highly confident in the HMD (confidence score of 10), their confidence was better-grounded in the recall accuracy, than when on the desktop.

4.3 Errors and Skips

The recall accuracy measures the number of correct answers. The participants in our user studies made an error in recall (i.e. gave an incorrect answer) or skipped answering (i.e. did not provide an answer). We show the percentile distribution of the average number of erroneous answers per participant for each display modality in Figure 7. Participants in a HMD made on average fewer errors than those on the desktop. The total number of errors on the head mounted display for 40 people were 33 out of 840, and on the desktop was 56 also out of 840. In addition, the difference in the incorrect answers is statistically significant, with $p = 0.0463 < 0.05$.

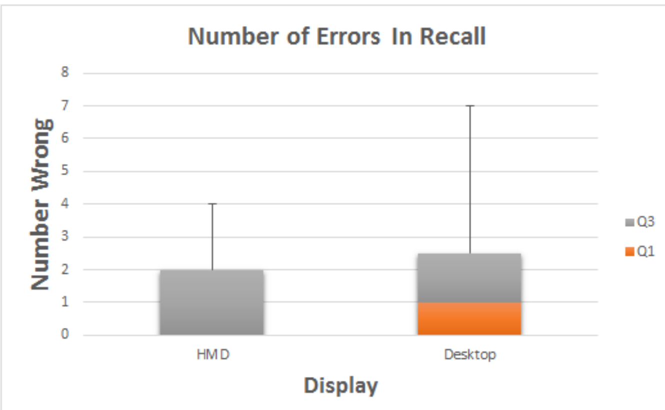

Figure 7: The average number of incorrect answers for each display modality. Note that the median number of incorrect answers for the HMD is 0 and therefore the first quartile is also 0.

In Figure 8, we show that the number of faces for which the participants skipped an answer on the desktop were significantly higher than those in the HMD. This was statistically significant with $p = 0.0062 < 0.05$, which reinforces that the participants in the HMD had better recall than those on the desktop.

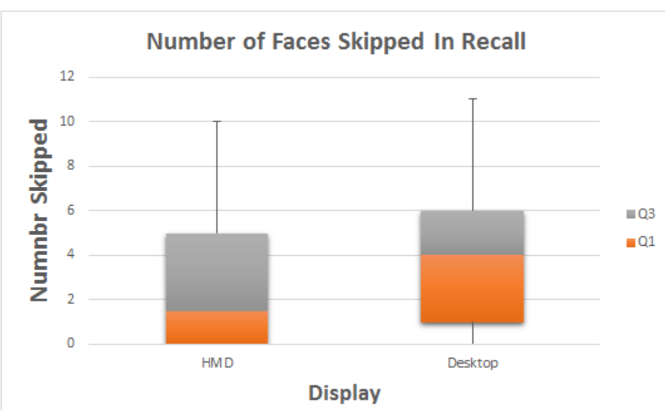

Figure 8: The number of faces skipped during recall for each display modality.

4.4 Recall Time

We also recorded the time it took for participants to recall. This was not designed to draw any definitive conclusions, but rather to be more complete in our data collection. Each user was given five minutes to train and up to five minutes for evaluation. The Figure 9 below shows the distribution of the average test times for the two display modalities. We did not ask the participants to finish as early as possible. As a result of this we noted a wide variation in the participant behavior. If a participant was confident in their answers they would end early. However, sometimes a participant would have answered everything, but would prefer to double check all their answers. Analogously, participants that did not know the answers to all the questions would end early because they gave up or would spend the full five minutes trying to recall the missing answers. Time therefore did not directly correlate with accuracy or how well a participant remembered, but simply to their style or confidence. Not surprisingly, we found that the difference in the testing times between the two modalities was not statistically significant.

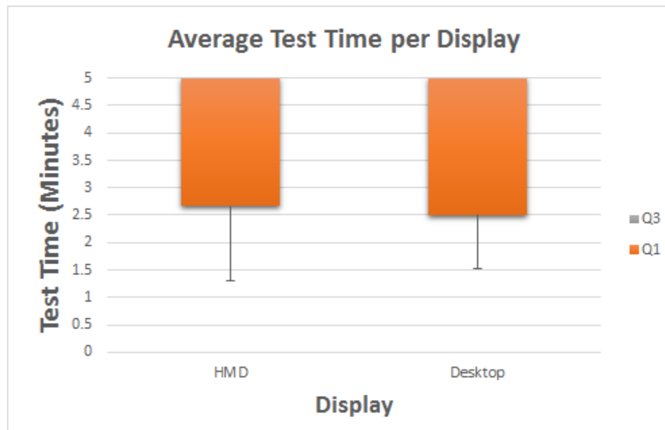

Figure 9: The average amount of time it took participants to recall all the faces in the HMD and on the desktop.

4.5 Alternating Displays

We show that there is an improvement in using a HMD compared to a desktop for a single user. Figure 10 shows the accuracy when using a desktop first then a HMD and using a HMD first and then a desktop.

For both the desktop and HMD, users started with roughly the same performance on both the desktop and HMD. However, when going to the other display, the performance changed. When users went from a desktop to a HMD, their performance generally improved. However, when the users went from a HMD to a desktop, their performance surprisingly decreased. ANOVA found that there is no statistically significant difference between the ordering of displays, with $F(1, 78) = 1.7$ and $p = 0.19$.

5 Discussion, Conclusions, and Future Work

We have found that the use of virtual memory palaces in HMDs is better than traditional desktops when trying to memorize and recall information. We had 40 participants memorize and recall faces on two display modalities for two virtual memory palaces, with two different sets of faces. Our study participants went through each permutation of those sets. In the end, there was roughly a 10% difference in

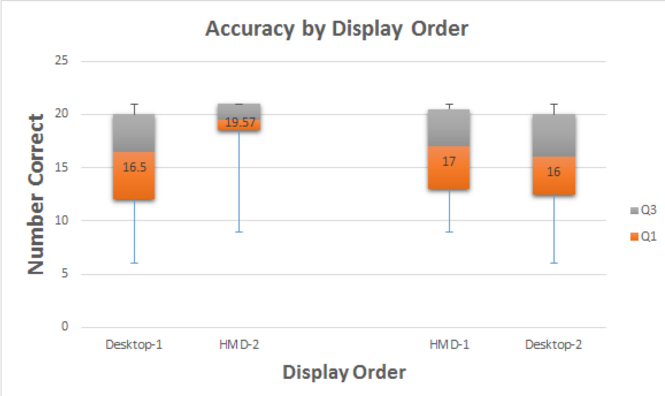

Figure 10: The average performance of participants going from a desktop to a HMD and from a HMD to a desktop.

the average recall accuracy between the two displays, which was found to be statistically significant. Virtual reality can clearly provide entertaining and immersive experiences, but we wanted to explore ways in which VR could enhance productivity. Given the results of our user study, we believe that virtual memory palaces offer us an exciting glimpse into how we may be able to organize and structure large information spaces and navigate them in ways that assist in superior recall.

During the user study, we offered the participants a chance to give us comments and also recorded anything interesting we observed. Nearly 75% of the participants pointed at objects while they were memorizing and recalling in HMDs. This behavior was far less common on the desktop. All of our study's participants were expert desktop users, but almost none had experienced a HMD before. We believe that if there were to be any implicit advantage it would lie with the desktop, given the overall familiarity with it. Although we gave the participants enough time to get comfortable in the HMD, before we began the study but we observed that many were still distracted and looked around while still performing better in the headset compared to the desktop. Lastly, we asked each participant which display they preferred in terms of achieving the task of recall. We explicitly stated that their decision should not be based on the novelty or "coolness" of the display or the experience. All but 2 of the 40 participants stated they preferred the HMD for this task. They further stated that they felt more immersed in the scene and so could be more focused on the task. In addition, they reported that HMD afforded them a superior sense of the spatial awareness which they claimed was important to some of their success. Several mentioned that they exploited the virtual memory palace setup and associated information relative to their own body.

Our study provides a tantalizing glimpse into what may lie ahead in virtual-environment-based tools to enhance human memory. The next steps will be to identify and characterize what elements of virtual memory palaces are most effective in eliciting a superior information recall. This could include elements in the architecture of the virtual memory palaces such as their design, their type, and various kinds of layouts and distribution of content that could help with recall. Another interesting future work would be to allow people to build their own virtual memory palaces, manipulate and organize the content on their own, and then ask them to recall that information. If their active participation in the organization of the data in virtual memory palaces makes a meaningful difference, then that could be further useful in designing interaction-based virtual environments that could one day assist in far superior information management and recall tools than those currently available to us. Yet another interesting future direction of research could be to explore elements of virtual memory palaces that are highly personal and those that could be used by larger groups. Much as textbooks and videos are used today for knowledge dissemination, it could be possible for virtual memory palaces to be one day used for effective transfer of mnemonic devices amongst humans in virtual environments.

Acknowledgements

We would like to thank Barbara Brawn-Cinani and Tom Ventisas for their support with our research. In addition we would like to thank the graduate students of the GVIL lab at the University of Maryland for participating in our pilot study and for their continued support. We appreciate the support of the NSF grant 14-29404, the State of Maryland's MPower initiative, and the NVIDIA CUDA Center of Excellence. Any opinions, findings, conclusions, or recommendations expressed in this article are those of the authors and do not necessarily reflect the views of the research sponsors. Lastly, we would like to thank the 40 participants who came to our lab and participated in our user study, without whom this research would not have been possible.

References

[1] D. A. Bowman and R. P. McMahan. Virtual reality: how much immersion is enough? *Computer*, 40(7):36–43, 2007.

[2] B. M. Brooks. The specificity of memory enhancement during interaction with a virtual environment. *Memory*, 7(1):65–78, 1999.

[3] A. Cockburn. Revisiting 2d vs 3d implications on spatial memory. In *Proceedings of the Fifth Conference on Australasian User Interface - Volume 28*, AUIC '04, pages 25–31, Darlinghurst, Australia, Australia, 2004. Australian Computer Society, Inc.

[4] J. Harman. Creating a memory palace using a computer. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '01, pages 407–408, New York, NY, USA, 2001. ACM.

[5] E. A. Johnson and N. Adamo-Villani. A study of the effects of immersion on short-term spatial memory. *World Academy of Science, Engineering and Technology*, 71:582–587, 2010.

[6] E. L. Legge, C. R. Madan, E. T. Ng, and J. B. Caplan. Building a memory palace in minutes: Equivalent memory performance using virtual versus conventional environments with the method of loci. *Acta Psychologica*, 141(3):380 – 390, 2012.

[7] K. Mania and A. Chalmers. The effects of levels of immersion on memory and presence in virtual environments: A reality centered approach. *CyberPsychology & Behavior*, 4(2):247–264, 2001.

[8] K. Mania and C. Randell. Monitoring navigational strategies and idle time in real and virtual environments: an experimental study. In *Proc. of the 8th International Conference on Virtual Systems and Multimedia 2003 (VSMM)*, pages 327–335, 2002.

[9] K. Mania, T. Troscianko, R. Hawkes, and A. Chalmers. Fidelity metrics for virtual environment simulations based on spatial memory awareness states. *Presence: Teleoperators and Virtual Environments*, 12(3):296–310, 2003.

[10] A. McAfee and E. Brynjolfsson. Big data: the management revolution. *Harvard business review*, (90):60–6, 2012.

[11] R. Pausch, D. Proffitt, and G. Williams. Quantifying immersion in virtual reality. In *Proceedings of the 24th Annual*

Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '97, pages 13–18, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.

[12] E. D. Ragan, A. Sowndararajan, R. Kopper, and D. A. Bowman. The effects of higher levels of immersion on procedure memorization performance and implications for educational virtual environments. *Presence: Teleoperators and Virtual Environments*, 19(6):527–543, 2010.

[13] B. E. Riecke, B. Bodenheimer, T. P. McNamara, B. Williams, P. Peng, and D. Feuereissen. Do we need to walk for effective virtual reality navigation? physical rotations alone may suffice. In *Spatial Cognition VII*, pages 234–247. Springer, 2010.

[14] G. Robertson, M. Czerwinski, K. Larson, D. C. Robbins, D. Thiel, and M. Van Dantzich. Data mountain: using spatial memory for document management. In *Proceedings of the 11th annual ACM symposium on User interface software and technology*, pages 153–162. ACM, 1998.

[15] A. Sowndararajan, R. Wang, and D. A. Bowman. Quantifying the benefits of immersion for procedural training. In *Proceedings of the 2008 Workshop on Immersive Projection Technologies/Emerging Display Technologiges*, IPT/EDT '08, pages 2:1–2:4, New York, NY, USA, 2008. ACM.

[16] M. Wraga, S. H. Creem-Regehr, and D. R. Proffitt. Spatial updating of virtual displays. *Memory & cognition*, 32(3):399–415, 2004.