

Elucidating Factors that can Facilitate Veridical Spatial Perception in Immersive Virtual Environments

Victoria Interrante¹, Brian Ries¹, Jason Lindquist¹, and Lee Anderson²

¹Department of Computer Science, University of Minnesota

²Department of Architecture, University of Minnesota

ABSTRACT

Enabling veridical spatial perception in immersive virtual environments (IVEs) is an important yet elusive goal, as even the factors implicated in the often-reported phenomenon of apparent distance compression in HMD-based IVEs have yet to be satisfactorily elucidated. In recent experiments [e.g. 3], we have found that participants appear less prone to significantly underestimate egocentric distances in HMD-based IVEs, relative to in the real world, in the special case that they unambiguously know, through first-hand observation, that the presented virtual environment is a high fidelity 3D model of their concurrently occupied real environment. We had hypothesized that this increased veridicality might be due to participants having a stronger sensation of ‘presence’ in the IVE under these conditions of co-location, which state of mind leads them to act on their visual input in the IVE similarly as they would in the real world (the *presence hypothesis*). However, alternative hypotheses are also possible. Primary among these is the *visual calibration hypothesis*: participants could be relying on metric information gleaned from their exposure to the real environment to calibrate their judgments of sizes and distances in the matched virtual environment. It is important to disambiguate between the presence and visual calibration hypotheses because they suggest different directions for efforts to facilitate veridical distance perception in general (non-co-located) IVEs. In this paper, we present the results of an experiment that seeks novel insight into this question. Using a mixed within- and between-subjects design, we compare participants’ relative ability to accurately estimate egocentric distances in three different virtual environment models: one that is an identical match to the occupied real environment; one in which each of the walls in our virtual room model has been surreptitiously moved ~10% inward towards the center of the room; and one in which each of the walls has been surreptitiously moved ~10% outwards from the center of the room. If the visual calibration hypothesis holds, then we should expect to see a degradation in the accuracy of peoples’ distance judgments in the surreptitiously modified models, manifested as an underestimation of distances when the IVE is actually larger than the real room and as an overestimation of distances when the IVE is smaller. However, what we found is that distances were significantly underestimated in the virtual environment relative to in the real world in *each* of the surreptitiously modified room environments, while remaining reasonably accurate (consistent with our previous findings) in the case of the faithfully size-matched room environment. In a post-test survey, participants in each of the three room size conditions reported equivalent subjective levels of presence and did not indicate any overt awareness of the room size manipulation.

CR Categories and Subject Descriptors: I.3.6 [Computer Graphics]: Methodology and Techniques; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism.

Additional Keywords: egocentric distance perception, immersive virtual environments.

1. INTRODUCTION AND PREVIOUS WORK

Because virtual environments technology has the potential to enable the presentation of information in the context of a user’s natural, egocentric frame of reference, it has great promise as an enabling technology for immersive design and evaluation in fields such as architecture and engineering where designers and others can benefit from experiencing virtual models at true scale. However, when considering the practical use of this technology, it is important to have a reliable understanding of whether, under what conditions, and to what extent the user is likely to be able to actually achieve an accurate understanding of spatial relationships (i.e. size and distance) within the virtual environment.

Numerous previous studies [e.g. 2, 13] have looked at this question, and many [e.g. 12] have found that people appear to systematically underestimate egocentric distances in immersive virtual environments presented via head mounted display systems. Although real and virtual environments differ in many significant respects, the particular factors influencing the accuracy of peoples’ distance estimates in immersive virtual environments have yet to be clearly identified. Loomis and Knapp [7] provide an excellent review of the extensive literature on this question.

Over the past several years, we have undertaken a series of experiments intended to elucidate the factors influencing peoples’ judgments of egocentric distances in immersive virtual environments with the aim of gaining insight into potential methods for facilitating more accurate distance perception in these environments. Initially, we discovered that people seem to be able to make accurate judgments about egocentric distances in an immersive virtual environment when the IVE represents a high fidelity model of the same physical space that the user is actually occupying, and the user has been able to unambiguously verify this by viewing the real space prior to donning the display upon which the corresponding virtual environment is presented [3]. Recently, we have replicated these findings in situations where the IVE represents a high fidelity model of a real space that the user is not currently occupying but has recently spent time in [4].

One possible interpretation of these intriguing results is that observers are better able to make accurate judgments of egocentric distance in an immersive virtual environment when they are able to become cognitively immersed, or *present*, in the IVE – i.e. when they are able to accept the virtual environment as being equivalent to the real world and therefore to act on their visual input in the virtual world in the same way that they would in the real world. However, other interpretations are also possible: for example, it could be that when people are exposed to a virtual environment that exactly corresponds to a real environment that they have just seen, they are able to make accurate judgments of egocentric distances in that IVE because they were able to form a metrically accurate mental model of the spatial structure of the real environment upon their brief exposure to it, and when they are subsequently presented with the corresponding virtual environment they are able to calibrate their interpretation of sizes and distances in the visually presented IVE to be consistent with their remembered model of the real environment.

It is important to be able to differentiate between these two hypotheses because each suggests a different strategy for attempting to facilitate accurate distance perception in an immersive virtual environment that is *not* a faithful replica of an actual existing environment, which is our ultimate practical objective. For example, to the extent that it is a question of enhancing presence, we would want to consider taking steps such as tracking users' hands and feet and providing them with a faithful representation of their own body in the virtual environment [10], or giving them the opportunity to physically interact with real objects that are also tracked and represented in the virtual environment [6]. To the extent that it is a question of providing reliable indicators of familiar size, we might work on enhancing the virtual environment with *entourage elements*, such as models of people, as architects often use to provide a sense of scale in their drawings. The immersive modeling software that we have developed for use in our design studio classes [1] already includes this feature.

2. OUR EXPERIMENT

In order to disambiguate the presence hypothesis and the visual calibration hypothesis, we designed the following study. Using a mixed within- and between-subjects design, we asked observers to make judgments of egocentric distance in a real room and in one of three different virtual environment models, each of which was described, via written instructions, as representing a "high fidelity virtual model of that same room". However, only one of the virtual models was actually an identical match in size to the real room. One third of the participants viewed a virtual model in which each of the walls had been surreptitiously moved 10% inward towards the center of the room (and the textures touched up in Photoshop to hide this change, without scaling anything), and another third viewed a virtual model in which each of the walls had been surreptitiously moved 10% outwards from the center of the room (and the textures appropriately filled in to hide this change, without scaling anything). The explanation for this choice of experimental design is as follows. In our previous studies, we had anticipated, and noted, some subtle systematic individual differences in participants' distance estimations made using the blind walking metric, with some people walking a bit longer than average, and some walking a bit shorter than average, consistently across conditions – an observation that was later confirmed, in the case of real-world blind walking, in large retrospective study by [5]. Therefore we felt it important to design our current study so that each person could serve as his or her own control to the greatest extent possible. However, because we did not want to overtly inform participants about the room size manipulation, we felt it essential to have each participant experience only one of the three virtual room models. In addition, although in our previous studies we had found no differences between conditions in which participants made distance judgments in the real world first versus in the virtual world first [3], in this experiment we felt that it would be important to have participants always perform the trials in the virtual environment first, out of concern that extensive prior physical experience in the real room might increase the likelihood of their consciously noticing any size mis-match between the real and virtual rooms.

In sum, through this experiment we aim to see how subtle, covert manipulations in the size correspondence (smaller /same /larger) between real and virtual room models (that participants are led to believe correspond exactly) might affect the *difference* between participants' real-world and virtual-world distance judgments. In light of our previous finding that participants appear to estimate distances with approximately equivalent accuracy in our matched real and virtual room environments, if the visual calibration hypothesis holds then we should expect to

find that participants who see the *smaller* room will *overestimate* distances in the virtual environment relative to in the real room, and that participants who see the *larger* room will *underestimate* distances in the virtual environment relative to in the real room. For example, if the real room is 30' long, and the virtual room model is 24' long, but participants interpret it as being 30' long, then when they are asked to close their eyes and walk towards a marker that is placed at a distance halfway down the virtual room model, we would expect them to walk 15' in the direction of the marker, rather than 12'. Likewise, if the virtual room is 37.5' long, but participants interpret it as being 30' long (matching the size of the real room), then when they are asked to blind walk towards a marker that is placed halfway down the length of the virtual room model, we would expect them to walk only 15' in the direction of the marker rather than the full extent of the marker's true distance, 18.75'. On the other hand, if the presence hypothesis holds, then we should expect that participants will judge distances with equivalent accuracy under all three virtual room conditions if the size manipulation is completely unperceived, or, that they will make similar errors in each of the manipulated conditions, to the extent that they might subjectively perceive the modified virtual room environments as being in some way unreliable representations of the real room.

2.1. Method

As in our previous experiments, we used 'blind walking' [9] to assess distance perception. Although some questions have been raised about potential problems with this metric [e.g. 8], due to a lack of good alternatives it remains the most commonly-used metric for judgments of egocentric distances in virtual environments over intervals of less than 20m. We used written instructions to enforce consistency in the presentation of information and instructions to participants across groups. All participants were informed that they would be taking part in one of a number of experiments being undertaken as part of a larger study whose purpose was to "compare space and distance perception in virtual environments with space and distance perception in the real world under various different display and interaction conditions". They were further informed that that they would be participating in the condition "virtual room, real room". Participants were *not* informed about the existence of different room models; on the contrary, each participant was *explicitly* told that the virtual room model that s/he would be seeing was an *exact replica* of the real room. Participants did not go through any training prior to testing and no feedback was made available to any participant about his or her performance at any time.

2.1.1. Apparatus

Testing took place in the Digital Design Lab located on the first floor of Walter Library on the University of Minnesota campus. This lab includes a fully tracked open space and a large, rear projected, curved screen display. The dimensions of the open space of the lab are 30' long x 25' wide in the center, tapering down to 16.5' wide at the edges due to the curvature of the screen. Figure 1 shows a photograph of the real-world lab environment.

The virtual environment was presented using an nVisor SX head mounted display manufactured by nVis. This visor provides 1280x1024 resolution images to each eye with an ~60° diagonal monocular field of view (for an effective resolution of about 2.2 arc minutes of visual angle per pixel) and 100% stereo overlap. The head mounted display is connected via a 15' cable to a video controller box stationed on a wheeled cart. This allows ample cord length to reach any point in the open space of the lab. We use a HiBall 3000 optical ceiling tracker, manufactured by 3rd Tech, to obtain information about the position and orientation of the user at a rate of about 500 Hz. With this tracker, and our real-

time rendering software, which was run on a PC with a 2.83GHz Intel Xeon processor with 2.0Gb of RAM and a Quadro 4900 XGL graphics card, we were able to present our simple virtual room model to our participants with minimal latency.



Figure 1: A photograph of the real room environment, illustrating a typical position of a participant at the start of a trial, and showing the arrangement by which the experimenter assisted with the management of the cables.

The original high fidelity virtual model of our real room environment was geometrically defined to be an exact match, in which each of the surfaces (floor, ceiling, and walls) was texture mapped with a mosaic of high resolution photographs obtained from the real room. There was a small amount of furniture in the real lab, such as chairs, computers and computer desks, but these were not included in the virtual model. For this experiment, we constructed smaller and larger versions of the original virtual room model by applying an $\sim\pm 20\%$ non-uniform scaling to the original model in the horizontal plane about the center of the room. Specifically, the scaling factors for the smaller and larger rooms were defined so that the ratio of the length of the longest wall in the smaller room to the length of the longest wall in the default room would be the same as the ratio of the length of the longest wall in the default room to the length of the longest wall in the larger room. This had the effect of moving each of the walls in by exactly 3 feet ($\sim 10\%$) towards the center in the case of the smaller room model, and out by exactly 3.75' from the center in the case of the larger room model, without changing the height of the room. We felt that it was important to leave the vertical extent of the rooms unchanged between conditions because of the potential complications that could be introduced if participants adopted a different understanding of their eye height in the different conditions as a result of the ceiling seeming closer or farther away from them in the virtual model than in the real space.

To accommodate the smaller or larger extents of visible floor, wall and ceiling surfaces in the modified room models, we had to define new textures for these surfaces, based on the textures used in the original. To create these new textures, we did not use any rescaling. Instead, for the wall surfaces, we took the prominent features, such as the doors and panels, reflected highlights, etc., and uniformly repositioned them, adding or subtracting white space using Photoshop's clone tool and touching up the result to hide any seams. For the ceiling surface we added or subtracted panels, taking care to maintain consistency between the locations of the light fixtures and the locations of the reflected highlights on the walls, and for the floor surface we simply extended or truncated the default repeating

texture pattern. Figures 2-4 show screenshots of each of the three different virtual room models taken from approximately the same position, representing a typical starting location at the beginning of a trial. From these images, it is clear that only the spacing of the prominent features was adjusted, while their size and general relative layout remained fixed.

The black electrical outlet floor plates that can be seen in figure 1 were omitted from all three of the virtual models in this experiment. This is a change from the situation in our previous experiments, in which the matching virtual room model included the floor plates. We felt that in this experiment the floor plates needed to be removed because there was simply no good way to incorporate them into the re-sized virtual models. Because the floor plates are occasionally stepped on during the trials (though we try to set things up to minimize this occurrence), we felt that it could be misleading to explicitly display them at positions in the virtual models that were offset from their actual positions in the real world. However, since the participants focus intently on the floor when making their distance judgments, we felt that if we were to leave the floor plates in their original positions relative to the center of the room in the resized room models, we would risk introducing an obvious indication of a change in the wall positions between the real and virtual scenes, as the walls would begin to overlap the plates in the case of the smaller room model.

2.1.2. Participants

We recruited 23 participants for this study. None of these people had participated in any of our earlier experiments and all were naïve to the hypotheses underlying the current study. Nine of our participants were undergraduate students from various different departments at the University who were recruited through a filmmaking/special effects interest group on campus; the remaining fourteen were undergraduate students and teaching assistants from the Department of Architecture, recruited from a large design studio class. Participants' ages ranged from about 20 to 30, and they included 16 males and 7 females. Nine of the students experienced the smaller lab model, nine experienced the larger lab model, and five experienced the accurate (same sized) lab model, as a control and to verify consistency with our previous findings. Although we had initially planned to recruit participants only for the different-sized room conditions, and to rely on the results from our previous findings with the same-sized room model for comparison, we ultimately decided that it would be prudent to run additional, new participants in the same-sized room condition, in order to explicitly control for any possibility of effects due to any subtle differences in methodology between our current and prior experiments, such as the elimination of the floor plates or the recalibration of our tracking system subsequent to its re-installation in the room after having been moved to a different location for an intervening experiment. Each participant was given a \$10 gift certificate in compensation for his or her efforts.

2.1.3. Procedure

All participants began by entering the lab and sitting down at a desk to read the written instructions and sign the consent form. After this, they were given verbal instructions about how to put on and adjust the head mounted display for optimal viewing. They were then guided to the edge of the room, where they put on the head mounted display and a small portable radio with headphones, which was used to provide pink noise to drown out any possible subtle auditory cues from the surrounding lab environment. Participants viewed the virtual model from a stationary position with an example target location displayed



Figure 2: An image of the **smaller** virtual lab model, in which the walls were moved 10% inward toward the center of the room.



Figure 3: An image of the original (**same sized**) virtual lab model.



Figure 4: An image of the **larger** virtual lab model, in which the walls were moved 10% outward from the center of the room.

while the procedure for the blind walking task was described to them again, verbally. After this, they began the experiment.

We had each person perform 20 trials of blind walking in the virtual environment, followed by 10 trials of blind walking in the real world. The number of trials was chosen so that the amount of time participants spent in each condition would be approximately equal. Participants were allowed to re-acclimate to the real world before testing by taking a short break between the virtual- and real-world trials. Each trial consisted of a direct blind walk from the participant's current location to a target 'tape' mark positioned at a randomly determined distance 8-25 feet away from the participant along their direction of view. We used a boundary condition to ensure that all of the target locations generated by the random process were displaced by at least 8 feet from both the virtual and real walls in all conditions. All procedures were conducted identically, regardless of which room model the participant experienced.

Two people were involved in running the experiment. During the virtual-world trials, one person (the operator) ran the keyboard controls at the computer, while the other person (the assistant) managed the cables for the participant, both keeping them out of his way and relieving any backwards tugging on the headset due to their weight. To simulate a blindfold in the virtual environment, the images to the head mounted display were cleared to black while the participant was walking. For timing purposes the participants announced when they were ready to close their eyes and begin walking and at that signal the display was blacked out by the operator. When the participant felt he had reached the target tape's location, he stopped walking and announced that he was done, and the operator used this as his signal to record the ending position. To prevent the participant from gaining any insight into the accuracy of his performance, the assistant then gave him verbal commands to walk in a circuitous route to a different location while he kept his eyes shut and the display remained turned off. Because the tape marks were placed virtually, the assistant managing the cables was generally unaware of their corresponding location in the real room, and it is highly unlikely that he would be able to subconsciously influence the participant to walk shorter or longer on any trial.

Distance interval endpoints in the real-world trials were indicated by two thin strips of cloth, sewed to pieces of Velcro, which were applied to the floor at random locations by one of the experimenters just before the beginning of each trial, and out of the sight of the participant. Participants began each trial by lining up with one of these 'tape' marks, taking visual aim at the other, then putting on a blindfold, closing their eyes, and walking. When the participant reached their estimated target location, the two experimenters used a tape measure to record the distance of the walk and the distance between the cloth markers. The participant was then verbally instructed to move to a different starting location while remaining blindfolded, and the strips were repositioned for the next trial.

Upon completion of all walking trials, the participants were seated once more and asked to fill out a two-page questionnaire regarding their experience. Although it has been shown that questionnaires are generally not a reliable tool for determining the extent to which a person feels 'present' in a virtual environment [11], our primary intent was to use the questionnaire not so much to assess presence as to provide a device for encouraging participants to let us know if they noticed anything 'not right' about the virtual environment. On the first page, participants were asked to provide ratings, on a scale from 1 to 7, about




Figure 5: A scatter plot showing all distance judgments made by the nine participants who viewed the **smaller** virtual room, color coded by participant ID. Squares indicate judgments made in the real world; circles represent judgments made in the virtual world.




Figure 7: A scatter plot showing all distance judgments made by the nine participants who viewed the **larger** virtual room, color coded by participant ID. Squares indicate judgments made in the real world; circles represent judgments made in the virtual world.




Figure 6: A scatter plot showing all distance judgments made by the five participants in our current experiment who viewed the **same-size** virtual room. Squares indicate judgments made in the real world; circles represent judgments made in the virtual world.

various aspects of their experience in the virtual environment. On the second page they were asked to provide an open-ended response to the following question: “Please describe in detail each of the characteristics of the presented virtual environment, or your experience in it, that felt unnatural or that you think might have detracted from your ability to function in the virtual environment in the same way that you would have functioned in the real world. We appreciate all of the insights that you can offer.”

2.2. Results

Figures 5 through 12 illustrate the results from this experiment. Figures 5-7 show scatter plots of each of the individual distance judgments made by each of the participants in each of the three




Figure 8: A point plot showing the average relative error in distance judgments made in the real vs. virtual environments by participants experiencing the **same-size** virtual room condition. Data from the current participant cohort is shown in dark brown; data from our previous participant cohort is shown in light brown.

room size conditions. Figures 8-10 show individual point plots of the average relative errors in distance judgments made by each participant in the real versus virtual environments under the three different room size conditions. For comparison purposes, figure 8 also shows data from the five participants in our previous experiment [3] who performed the identical task under the matched size condition. Points are rendered as solid when the difference between a participant’s performance in the real and virtual worlds was found to be strongly statistically significant ($p < 0.01$) and rendered with a small white dot in the center when the difference was significant with $p < 0.05$. Points are rendered as hollow (with a large white dot in the center) when the difference between a participant’s performance in the real and virtual




Figure 9: A point plot showing the average relative error in distance judgments made in the real and virtual environments by participants experiencing the **smaller**-size virtual room condition.

environments was not significant ($p \geq 0.1$). We can see that, as before, most participants who experienced the same-sized virtual room model judged distances with similar accuracy in both the real and virtual environments. However, in figures 9 and 10, we can see that most of the participants who experienced the smaller and larger room models *underestimated* distances, on average, to a greater extent in the virtual world than in the real world.

To verify these observations, we performed a statistical analysis, via ANOVA, on the effect of technology (trials done in the real world vs. trials done in the virtual world) on the magnitude of the average relative errors observed in each of the differently sized room conditions. (While of course it is not necessary to use an ANOVA to examine the effects of a single condition, there is no reason to expect that an ANOVA would give different or less accurate results than a simpler alternative method.) In the case of the same-sized room data, pooling the data from the ten total participants who experienced this condition in our current and previous experiments, we found no significant main effect of technology (real world vs. virtual world) on errors in distance judgments $\{F(1,18) = 0.7376, p = 0.402\}$. Looking only at the data from the five participants in our present study, the result is basically the same $\{F(1,6) = 0.3253, p = 0.584\}$. However in the case of the larger-sized room, we did find a significant main effect of technology $\{F(1,16) = 14.07, p = 0.0017\}$, and the same was true in the case of the smaller sized room $\{F(1,16) = 8.581, p = 0.0098\}$.

Of course, the fact of our finding no significant effect of technology in the case of the same-sized real and virtual rooms is *not equivalent* to finding that distance estimation accuracy is the *same* in these two cases, since one cannot use statistical analysis to prove the null hypothesis. Nevertheless, the fact that we see that distances are significantly underestimated in the virtual environment relative to in the real environment when the virtual and real rooms are *differently sized* and *not* when they are the *same* size suggests a significant effect of the room size manipulation. To verify this, we ran an additional ANOVA analysis to directly test whether there is a significant effect of the

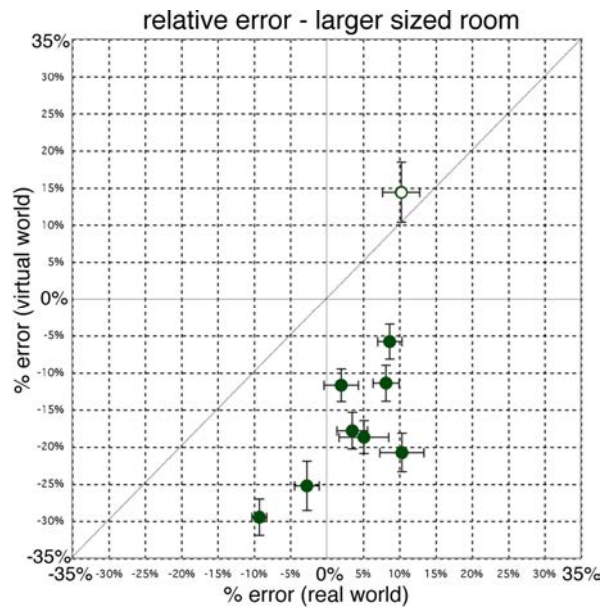


Figure 10: A point plot showing the average relative error in distance judgments made in the real and virtual environments by participants experiencing the **larger**-size virtual room condition.

room size manipulation condition on the magnitude of the difference between the average relative errors that participants make in the virtual room relative to in the real room. In this analysis, we found a significant main effect of the room size condition $\{F(2,25) = 8.716, p = 0.001343\}$. Running a Tukey HSD test on the pairwise differences between errors in the three room conditions, we found that the underestimation of distances in the virtual world relative to in the real world was significantly greater ($p < 0.05$) in the larger virtual room condition than in the same-sized virtual room condition, and marginally significantly greater ($p < 0.1$) in the smaller virtual room condition than in the same-sized virtual room condition.

Unfortunately, we were not able to gain much useful information from participants' responses to the questions on our 'presence' questionnaire, as the (between-subject) responses were generally similar across all three virtual room conditions. The only significant differences we found were in the responses to questions 1 and 2 in the cases of the smaller versus same-sized environments. Question 1 asked "How 'real' did the depicted virtual environment *look* to you, while you were in it, on average?" and question 2 asked "How 'real' did the depicted virtual environment *feel* to you, while you were in it, on average?" The mean response to these questions from the 9 participants who experienced the smaller room was slightly lower than the mean response from the 5 participants who experienced the same-sized room in this experiment. The other questions asked: 3) "How comfortable did you feel in the presented virtual environment?", 4) "To what extent did you feel that your ability to perform actions in the presented virtual room resembled your ability to perform these same actions in the corresponding real room?", 5) "To what extent, on average, did you feel as if you were actually physically present in the environment depicted by the head mounted display, over the course of the experiment?", 6) "To what extent/how often did you think about the fact that the room presented via the head mounted display represented the same physical space as the room you were actually in?", and 7) "To what extent/how often did you feel as if the room presented via the head-mounted display didn't 'feel the same' as the room you were actually in?" Figure 11 shows a bar graph of the questionnaire results.




Figure 11: Presence questionnaire responses, averaged over all participants in each of the different room conditions. The error bars show the 95% confidence intervals of the means in each case.

In their responses to the final, open-ended question, only two of the 23 participants mentioned anything about the size of the room seeming ‘off’. One participant in the larger room condition remarked that “the virtual room was empty which makes it feel somewhat bigger”, and another participant in this same condition wrote that “the room looked smaller, and it looked like the ground was closer than normal”. In addition, one person in the smaller room condition remarked that she “somehow felt taller with the virtual reality”, which is consistent with an underestimation of distances. However to the majority of participants it appeared as if the virtual and real rooms were perceived to be good matches, in terms of size. One participant in the larger room condition wrote that he was “amazed by the head mounted display and how realistic in proportion and scale it was”. Similarly, a participant in the smaller room condition wrote that “the dimensions felt very similar to the actual room, which helped me in the virtual simulation because I had already seen the real room”. Nearly all of the participants commented on various other factors related to the difference between the virtual- and real-world experience but unrelated to the perception of size differences. These factors included, in rough order of frequency: the more limited field of view in the head mounted display; noticing a latency in the display, especially when they swung their head around; being disturbed by the inability to see their feet; not being able to see the furniture in the virtual environment, which for one person induced a concern about walking into unseen objects; being able to feel the floor sockets in the real room but not to see them in the virtual room; perceiving a ‘fuzziness’ in the image presented by the HMD; being “distracted by the weight of the head mounted display and the way the screen moved [a little shakey/wobbly]”; the virtual environment seeming brighter; everything looking “too perfectly crisp” in the HMD; and noticing that the moldings around the door were not being modeled in 3D. Finally, one person in the larger room condition wrote: “I feel that the environment of virtual reality has very much difference from reality, and this makes me feel unsecured to walk freely”.

2.3. Discussion

The main conclusion that we can draw from the results of this experiment is that the ‘visual calibration’ hypothesis is not very well supported by the data. In particular, if participants had been using information gleaned from their exposure to the real room to calibrate their perception of the size of the virtual room, we would have expected to see *opposite* effects on distance judgments as a result of expanding versus shrinking the virtual room model.

Instead, we found that distance perception accuracy was diminished in the *same* way (towards an underestimation of distances) in each of these cases. While these results do not provide strong direct support for the presence hypothesis, they are not inconsistent with what would be predicted by the presence hypothesis if manipulating the size of the virtual room away from being an exact match to the real room somehow caused participants to have greater doubts about the virtual room’s reliability as a faithful and exact representation of reality. Unfortunately our questionnaire responses do not provide strong support for this interpretation; however they also do not provide support for any alternative interpretation, though it may well be that different factors are affecting the accuracy of participants’ distance judgments in the different cases.

One interesting item of note is that in examining the data from our current experiment, in conjunction with related data from our previous experiments, we noticed a consistent trend for participants to ‘walk longer’ in successive trials in each of the virtual room conditions but not in the real world, suggesting that there may be some general effect on participants’ distance estimates of time spent in the virtual environment, at least when the blind walking metric is used. Figure 12 shows the data for the virtual-world trials arranged by trial number, and figure 13 shows the data for the real-world trials. In each case, the relative errors in participants’ distance judgments, $(\text{target_dist} - \text{walked_dist}) / (\text{target_dist})$, are separately averaged by trial number over all of the participants in each condition, and trend lines showing the best linear fit to the average data are superimposed where the trend was significant at $p < 0.05$, according to r_s and t values computed using a Spearman rank order correlation.

We also noticed a trend for the relative error in participants’ distance estimates to increase with the length of the distance interval traversed, becoming noticeably larger for distances estimates obtained over intervals of greater than 20 feet. This is consistent with there being a possible effect of the known finite size of our room on participants’ willingness to walk long distances without sight, though other explanations are of course also possible. However we would need to collect data from more participants under a wider range of conditions before we could either confidently assert the significance of this trend or speculate about its potential basis.

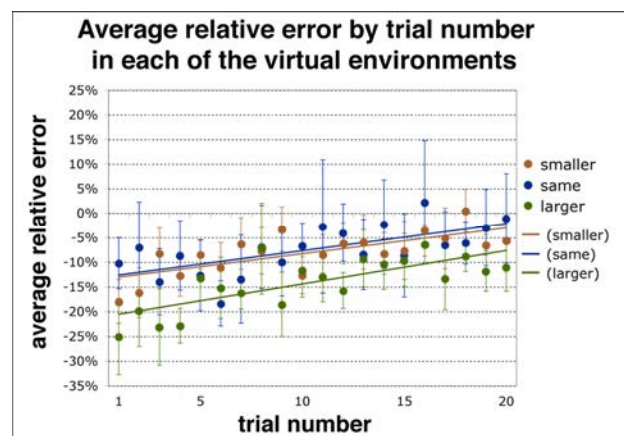


Figure 12: The average relative errors in participants’ distance estimates by trial number in each of the three different virtual room conditions.

3. FUTURE WORK

In future work, we would like to explore the development of robust, proactive strategies for facilitating more accurate distance




Figure 13: The average relative errors in participants' distance estimates by trial number in the real world, after exposure to one of the three virtual environment room conditions. Only one trendline is shown here because the relationship between walk number and relative error is significant *only* in the condition in which the participants experienced the **larger** virtual room environment before performing the real-world trials.

perception in non-co-located virtual environments. Although we are aware of strategies that rely on manipulations of the visual stimulus to counter the effects of apparent spatial compression, we are concerned that such strategies could backfire if it turns out that what everyone is interpreting as spatial compression turns out to be merely an artifact of peoples' inherent biases under the conditions of uncertainty that arise when they are reluctant to assume the equivalence of the presented virtual environment to the real world. In that case, the amount of 'spatial compression' experienced would be a moving target, and fixed compensatory manipulations that initially seem to help could eventually backfire over time. As an alternative to that sort of approach, we are interested in studying the effects of pursuing techniques that have been shown to enhance participants' subjective sense of presence in a virtual environment, such as providing them with a visually and/or behaviourally faithful representation of their body using an auxiliary tracking system, or enhancing the virtual environment with ambient spatialized sound sources. In addition, we would like to directly test whether heavily populating a virtual environment with landmark objects that can provide good familiar size cues might enhance participants' ability to accurately judge sizes and distances in the virtual world. The use of these 'entourage elements' is a standard technique employed in architectural practice to help viewers to assume an appropriate interpretation of scale in drawings. Finally, we plan to re-examine the effects of near-range, sighted, active experience in a realistic but unfamiliar virtual environment on a participant's subsequent ability to accurately judge spatial relationships over farther extents in that environment.

ACKNOWLEDGMENTS

This research was supported by a grant from the National Science Foundation (IIS-0313226), by the University of Minnesota through a Digital Technology Center seed grant, and by the Linda and Ted Johnson Digital Design Consortium Endowment and Lab Setup Funds. We are grateful to the many people who have helped us to think through questions related to this work, particularly including Bernd Frölich, and we are indebted to an anonymous reviewer of one of our earlier papers for suggesting the possibility that the participants in our earlier studies might be

acquiring metric information about the real environment when they see it for the first time upon entering the space, and using that information to calibrate the perceived size of the matching virtual environment. We are also grateful to Ed Swan for suggesting that we take a closer look at whether the magnitude of our participants' errors changes over time. Finally, we are grateful to Mike Kaeding for helping to recruit participants for the studies reported in this paper, and to all of our participants for their dedicated and conscientious efforts.

REFERENCES

- [1] ANDERSON, L., ESSER, J. AND INTERRANTE, V. 2003. A virtual environment for conceptual design in architecture. In *Proceedings of the 9th Eurographics Workshop on Virtual Environments/7th International Workshop on Immersive Projection Technology*, Zurich, Switzerland, 57–63.
- [2] HENRY, D. AND FURNESS, T. 1993. Spatial perception in virtual environments. In *Proceedings of the 1993 IEEE Virtual Reality Annual International Symposium*, 33–40.
- [3] INTERRANTE, V., ANDERSON, L. AND RIES, B. 2006a. Distance perception in immersive virtual environments, revisited. In *Proceedings of IEEE Virtual Reality 2006*, 3–10.
- [4] INTERRANTE, V., ANDERSON, L. AND RIES, B. 2006b. Further evidence for cognitive influences on egocentric distance perception in immersive virtual environments, *manuscript in revision*.
- [5] KUHLMAN, S.A., CREEM-REGHEER, S.H., & THOMPSON, W.B. 2006. Individual differences in accuracy of blind walking to targets on the floor [Abstract]. *Journal of Vision*, **6**, 6, 726a, <http://journalofvision.org/6/6/726/>, doi:10.1167/6.6.726
- [6] LOK, B., NAIK, S., WHITTON, M.C. AND BROOKS, F.P., JR. 2003. Effects of handling real objects and self-avatar fidelity on cognitive task performance and sense of presence in virtual environments, *Presence: Teleoperators and Virtual Environments*, **12**, 6, 615–628.
- [7] LOOMIS, J.M. AND KNAPP, J.M. 2003. Visual perception of egocentric distance in real and virtual environments. In *Virtual and Adaptive Environments*, L.J. Hettinger and M.W. Haas, eds., 21–45.
- [8] PHILBECK, J.W. 2005. Rapid recalibration of locomotion during non-visual walking [Abstract]. *Journal of Vision*, **5**, 8, 308a, <http://journalofvision.org/5/8/308/>, doi:10.1167/5.8.308.
- [9] RIESER, J.J., ASHMEAD, D.H., TAYLOR, C.R. AND YOUNGQUIST, G.A. 1990. Visual perception and the guidance of locomotion without vision to previously seen targets. *Perception*, **19**, 675–689.
- [10] SLATER, M., USOH, M., AND STEED, A. 1994. Depth of presence in virtual environments. *Presence: Teleoperators and Virtual Environments*, **3**(2), 130–144.
- [11] SLATER, M. 2004. How Colorful was your Day? Why questionnaires cannot assess presence in virtual environments. *Presence: Teleoperators and Virtual Environments*, **13**, 4, 484–493.
- [12] THOMPSON, W.B., WILLEMSEN, P., GOOCH, A.A., CREEM-REGHEER, S.H., LOOMIS, J.M. AND BEALL, A.C. 2004. Does the quality of computer graphics matter when judging distances in visually immersive environments? *Presence: Teleoperators and Virtual Environments*, **13**, 5, 560–571.
- [13] WALLER, D. 1999. Factors Affecting the Perception of Interobject Distances in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, **8**, 6, 657–670.