

Perception of Spatial Relationships in Impossible Spaces

Andrew Robb
School of Computing
Clemson University
arobb@clemson.edu

Catherine Barwulor
School of Computing
Clemson University
cbarwul@clemson.edu

ABSTRACT

Impossible spaces have been used to increase the amount of virtual space available for real walking within a constrained physical space. In this technique, multiple virtual rooms are allowed to occupy overlapping portions of the physical space, in a way which is not possible in real Euclidean space. Prior work has explored detection thresholds for impossible spaces, however very little work has considered other aspects of how impossible spaces alter participants' perception of spatial relationships within virtual environments. In this paper, we present a within-subjects study ($n = 30$) investigating how impossible spaces altered participants' perceptions of the location of objects placed in different rooms. Participants explored three layouts with varying amounts of overlap between rooms and then pointed in the direction of various objects they had been tasked to locate. Significantly more error was observed when pointing at objects in overlapping spaces as compared to the non-overlapping layout. Further analysis suggests that participants pointed towards where objects would be located in the non-overlapping layout, regardless of how much overlap was present. This suggests that, when participants are not aware that any manipulation is present, they automatically adapt their representation of the spaces based on judgments of relative size and visible constraints on the size of the whole system.

CCS CONCEPTS

- Human-centered computing → User studies; Virtual reality;

KEYWORDS

virtual reality, impossible spaces, spatial perception, user studies

ACM Reference format:

Andrew Robb and Catherine Barwulor. 2017. Perception of Spatial Relationships in Impossible Spaces. In *Proceedings of Conference Name, Conference Location, Conference Date and Year*, 5 pages.

<https://doi.org/10.1145/8888888.7777777>

1 INTRODUCTION

Many different techniques have been developed to enable users to move through virtual spaces that are larger than the physical space available to them. Some of these techniques, such as teleportation [Bozgeyikli et al. 2016] or walking-in-place [Usoh et al. 1999], make

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Conference Name, Conference Date and Year, Conference Location

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-1234-5/17/07.

<https://doi.org/10.1145/8888888.7777777>

it obvious to the user that their motion through the virtual world is disjoint from their motion in the physical world. Other techniques, such as redirected walking [Nilsson et al. 2018], change blindness [Suma et al. 2011a], or impossible spaces [Suma et al. 2012], make use of users' perceptual limitations to mask this disconnection between the physical and the virtual space.

Redirected walking introduces subtle amplifications to a user's head rotation in order to influence a user's path through physical space. By manipulating rotation, users can be induced to walk in a curved path in the real world while appearing to follow a straight path in virtual reality (VR). This allows users to physically walk while exploring virtual spaces that are larger than the available physical space. A significant amount of work has explored perceptual thresholds for rotational gain [Steinicke et al. 2010] and translational gain [Grechkin et al. 2016], as well as how these thresholds can change over time [Bölling et al. 2019]. Other work has considered how redirected walking affects user's perception of spatial orientation within a virtual space [Langbehn et al. 2018] and cognitive load [Bruder et al. 2015]. Unlike redirected walking, change blindness allows users to explore larger virtual environments by dynamically changing the environment while users are distracted (e.g. changing where a door is located); Suma et al. found that change blindness did not impact people's memory of locations, as measured via a pointing task [Suma et al. 2011a].

Like redirected walking and change blindness, impossible spaces also enable users to physically walk through virtual spaces that are larger than the available physical space. However, instead of decoupling users' walking trajectory, impossible spaces instead allow different virtual rooms to occupy overlapping portions of the physical space; this is masked from users through the use of corridors to connect the overlapping rooms, so that users can never see that both rooms occupy the same physical space [Suma et al. 2012]. As with redirected walking, several studies have explored users' ability to perceive the impossible space at varying percentages of overlapping geometry [Suma et al. 2012], and how this is influenced by the complexity of the corridors connecting the overlapping rooms [Imura et al. 2015]. However, few studies have explored how impossible spaces may affect other aspects of spatial perception; Suma et al. found that users perceived distances between objects to increase as spaces overlapped more, as judged by a blind walking task [Suma et al. 2012].

In this paper, we present an initial exploration of how users' spatial knowledge is affected by impossible spaces. Spatial knowledge in VR has been investigated using a range of different techniques, including pointing towards remembered locations [Chance et al. 1998; Peck et al. 2011; Riecke and Wiener 2007; Ruddle and Lessels 2009; Suma et al. 2011b] and drawing maps of the space after exploration [Billinghurst and Weghorst 1995; Ruddle and Lessels 2006;

Sigurdarson et al. 2012; Suma et al. 2009; Zambaka et al. 2005]. Langbehn et al. used both techniques to assess the effect of redirected walking on spatial understanding, as compared to teleportation and joystick-based motion [Langbehn et al. 2018]; he found that participants ability to point at targets and draw a map of the environment was largely unaffected by locomotion technique, but that participants were better able to recall where furniture was placed in the room. However, it should be noted that the environment used in this experiment was a single large open room; this space is well suited for redirected walking, but poorly suited for impossible spaces, as it would be easy to detect the presence of overlapping geometry without walls separating the overlapping rooms.

We used similar methods in our experiment. Participants explored three different virtual environments containing different levels of overlap. After each trial, they pointed towards several objects placed in the virtual environment and then drew a map of the space they had just explored. Unlike Langbehn et al's findings, we observed that participants pointing error increased significantly ($p < 0.001$) in overlapping spaces. However, participants struggled to draw accurate maps in our experiment, possibly due to difficulties with the system we designed to allow them to draw maps in VR; we attempted to have participants draw maps in VR so as to avoid users needing to exit the VR environment in between trials. As participants struggled to draw useful maps, we focus on the pointing data in this paper.

2 METHODS

To investigate how a user's sense of spatial perception is affected by the degree to which spaces overlap, and by the dimensions of the space, we designed a within-subjects study where participants explored three different layouts (see Figure 1), each of which contained four rooms connected by hallways. Three variations of each layout were created: "normal", "expanded", and "compressed". In the normal version, the layout completely filled the available 6x7.5 meter physical space, and there was no overlap between rooms. In the expanded version, the same corridor design was used, but each room was expanded by 50% and overlapped other spaces; the shape of the rooms were not changed, and they were expanded in the direction that was deemed least likely to be noticed by participants. In the compressed version, the layout was shrunk to occupy a 5x6 meter space, and the rooms were created so that they occupied the same area as the equivalent rooms in the normal condition; this resulted in a similar overlapping percentage as was used in the expanded condition, but with less space dedicated to corridors.

2.1 Participants

Thirty participants were recruited for the study (22 males). Participants were an average of 24.0 ± 3.59 years of age, and reported having spent an average of roughly 15 hours in VR prior to the experiment, ranging from 0 to 50+ hours; participants were recruited from a VR course, however impossible spaces were not included in the curriculum. All participants were naive to the purpose of the experiment. The entire experiment took approximately 45 minutes.

2.2 Apparatus

A 6x7.5 meter tracked space was used during this experiment. Participants experienced the virtual environment using a HTC Vive Pro, coupled with the Intel Wireless expansion. Four Vive 2.0 Bases stations were used to track the entire volume. The simulation was run on a computer capable of maintaining a frame rate of 90 FPS for the duration of the experiment.

The virtual environment was designed using the Unity game engine. To create seamless transitions between the overlapping geometry, a set of triggers were carefully placed in each level that would swap out room geometry when users passed through them. These triggers were placed to ensure that participants never saw geometry that should not be visible based on their current progression through the space. The virtual environment was themed to resemble a generic office (see Figure 2).

2.3 Study Design

A mixed-design was used in this study, where all participants completed three trials; participants saw each layout (which we label A, B, and C) and each overlap once. Due to the large number of potential combinations, it was not feasible to counterbalance the order in which layout and overlap were presented. Instead, the pairing between layout and overlap was selected randomly, as was the order in which they were presented.

Upon arriving in the lab, participants were asked to sign an informed consent form and to complete a brief demographics survey. They then put on the HTC Vive HMD and entered an introduction scene. Participants were instructed that they would be navigating three different room layouts, that they should look for 4 labeled objects in each layout, and that they should remember where these objects were located. The actual objects used in each trial were selected randomly from a pool of candidate objects, each of which were shown to participants in the introduction scene; to avoid confusion with other objects in the scene, a highlight shader was applied to each object and a text label floated above it. Once participants felt they could remember where each of these objects were located, they pushed a button located at a specific location in each layout (no objects were visible from this location). They were then asked to point to the four objects, as identified by name, in a randomly determined order. Pointing was done using the HTC Vive controller, and a ray was emitted from the controller to indicate the pointing direction. After indicating the direction for each object, participants advanced to the map drawing stage. Once this was completed, the proceeded to the next trial. Upon completion of all three trials, participants completed two surveys and a short debriefing. Participants took an average of 260 ± 230 seconds to complete each trial. The entire experiment took roughly 1 hour to complete. No breaks were provided during the experiment.

2.4 Metrics

The location of the controller, the direction the controller was facing, the location of the object being pointed to, and the angular pointing error in the XZ plane was recorded for each trial (the Y direction was not considered as height was not manipulated between layouts). Participants' location in the virtual environment was also sampled at 30 Hz.



Figure 2: A typical room from the virtual environment. The plant in the left corner is an example of the type of object participants were to remember the location of.

Participants also completed a short demographics survey at the beginning of the study, and surveys measuring simulator sickness [Robert Kennedy, Normal Lane, Kevin Berbaum 1993] and feelings of presence [Schubert 2002] (scores on this scale range from -3 to 3) after all trials were completed. Surveys were not administered between each trial, as we did not expect the experimental manipulation to alter feelings of presence or sickness. Finally, participants completed a short debriefing and were asked what they thought the purpose of the experiment was. Only 2 participants reported suspicion that overlapping spaces were present in the study.

Participants reported moderate to high levels of presence (*Overall* = 2.29 ± 0.74 , *SpatialPresence* = 2.06 ± 0.71 , *Involvement* = 1.06 ± 1.07 , *PerceivedRealism* = 0.48 ± 0.79) and low to moderate levels of sickness (*Nausea* = 21.54 ± 25.24 , *Oculomotor* = 25.92 ± 25.50 , *Disorientation* = 27.84 ± 32.94 , *Overall* = 28.71 ± 29.15).




Figure 1: The Normal version of the three room layouts are shown above. The star represents the location participants stood when pointing, the circles represent the location of the objects pointed at, and the dark grey squares represent furniture (desks, chairs, or shelves). All doorways had doors which could be swung open, and which closed automatically when released. The arrows indicate the direction in which rooms were expanded. Furniture and objects in the room were also moved in the direction of the arrows when expanded. All sizes are approximately correct.

3 RESULTS

3.1 Pointing Error

As the design was not fully-crossed, each participant did not have data for all combinations of layout and overlap. As such, a mixed linear model was used to analyze the results. Room layout and overlap type were used as fixed effects, and participant ID was used as a random effect. P-values were obtained by likelihood ratio tests comparing the best-fit model to the model without the effect in question [Winter 2013].

For our analysis, we consider the absolute error made when pointing at a specified object. As the distribution of the absolute




Figure 3: Participants pointed more accurately in the Normal condition, as compared against the Compressed and Expanded condition.

error was non-normal, the data was transformed by taking the square root so as to create a more normal distribution for analysis (figures shown in this paper report the original data). The data was then mean centered. Outliers were removed using the IQR rule, with the threshold set to 1.5; this excluded all data points with error greater than 94.9 degrees, which accounted for 6.6% of the data. A visual analysis of these data points suggested that the large error was due to pointing at the wrong object, rather than misremembering exactly where an object was located.

A main effect was observed for overlap type ($\chi^2(1) = 20.503, p < 0.001$). A main effect was also observed for room layout ($\chi^2(1) = 25.154, p = 0.001$); this can most likely be attributed to the difference in pointing position between Layouts A and C and Layout B. No interaction effect was observed between overlap and room layout. Post-hoc pairwise comparisons revealed that participants were significantly more accurate in the Normal overlap condition, as compared with both the Compressed ($p = 0.002$) and the Expanded ($p < 0.001$) conditions; there were no significant differences between the Compressed and Expanded conditions ($p = 0.6305$). The pointing error for each of the three overlap conditions can be seen in Figure 3.

3.2 Secondary Analysis of Pointing Error

A visual analysis of the pointing error suggested that, in both the Compressed and Expanded conditions, participants pointed towards where the objects would be located if there was no overlap present (i.e. where the objects were located in the Normal condition). To test this, we determined which objects were moved in the Compressed and Expanded conditions, and which objects were not (these were generally placed in a corridor, e.g. the upper right object in Layout A). We then determined the point these objects would have been located in if no overlap had been present (this was not always identical to the location of the corresponding object in the normal condition, as corridors shifted slightly between overlap conditions). The angle between this corrected position and participants' pointing was then calculated. This yielded three different types of error: 1) error made when pointing at an object whose location never moved (these objects were either in the Normal overlap condition, or were placed in a corridor and thus were not affected by changes in room sizes), 2) the original error made when pointing at an object whose location was moved due to overlap, and 3) the corrected error made when pointing at an object whose location was moved due to overlap. The data for these three error types is shown in Figure 4.

We analyzed this data using a mixed linear model, as described above, with a fixed effect of error type and a random effect of participant ID. A main effect was observed for error type ($\chi^2(1) = 38.793, p < 0.001$). Post-hoc pairwise comparisons revealed that the Moved-Original error differed significantly from both Moved-Corrected ($p < 0.001$) and NotMoved-Original ($p < 0.001$), but that Moved-Corrected and NotMoved-Original did not differ significantly from each other ($p = 0.916$). This supports our observation that, in the Compressed and Expanded conditions, people pointed where the objects would have been located if no overlap was present.




Figure 4: Recomputing the error for moved objects based on a corrected position yields an error rate almost identical to the error observed for not moved objects.

4 CONCLUSION AND FUTURE WORK

Participants' error increased when pointing at objects in the impossible layouts, regardless of whether the space was Compressed or Expanded. Our secondary analysis suggested that participants pointed to where the objects would have been if the layout had not been Compressed or Expanded, which created the observed error. This suggests that, while making their judgments, participants relied on information that was not affected by the amount of overlap present in the given condition.

The objects participants pointed at were typically placed in a corner, or along a wall (see Figure 1). As such, participants may have mapped the location of these objects to the boundaries of the rooms they had explored. Given that all rooms were expanded by an equal proportion, it may be that participants recalled the relative sizes of the rooms, and then judged where the room boundaries were most likely to be based on the visible length of the hallway they were currently standing in. This approach would be especially relevant when standing in Layouts A and C, as participants made their judgments while standing in a corridor that bisected the environment. Future work could further investigate this hypothesis by testing configurations where some rooms were expanded while others were not, thus changing the relative proportions of the different rooms.

As participants were generally not aware that impossible spaces were present, it is also unclear how pointing accuracy would be affected if the impossible nature of the space was more apparent. Future work should investigate how people indicate an object's location in obviously impossible spaces, and if this is grounded in the relative position of the virtual spaces or in the absolute physical boundaries of the available space.

ACKNOWLEDGMENTS

The authors would like to thank Daniel Petty for his help in constructing the simulation used in this research. This work was funded in part by NSF Award 1717937.

REFERENCES

Mark Billinghurst and Suzanne Weghorst. 1995. The use of sketch maps to measure cognitive maps of virtual environments. In *Proceedings Virtual Reality Annual International Symposium'95*. IEEE, 40–47.

Luke Bölling, Niklas Stein, Frank Steinicke, and Markus Lappe. 2019. Shrinking Circles: Adaptation to Increased Curvature Gain in Redirected Walking. *IEEE transactions on visualization and computer graphics* (2019).

Evren Bozgeyikli, Andrew Raji, Srinivas Katkoori, and Rajiv Dubey. 2016. Point & teleport locomotion technique for virtual reality. In *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play*. ACM, 205–216.

Gerd Bruder, Paul Lubos, and Frank Steinicke. 2015. Cognitive resource demands of redirected walking. *IEEE transactions on visualization and computer graphics* 21, 4 (2015), 539–544.

Sarah S Chance, Florence Gaunet, Andrew C Beall, and Jack M Loomis. 1998. Locomotion mode affects the updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence* 7, 2 (1998), 168–178.

Timofey Grechkin, Jerald Thomas, Mahdi Azmandian, Mark Bolas, and Evan Suma. 2016. Revisiting detection thresholds for redirected walking: Combining translation and curvature gains. In *Proceedings of the ACM Symposium on Applied Perception*. ACM, 113–120.

M Imura, P Figueiroa, and B Mohler. 2015. Influence of path complexity on spatial overlap perception in virtual environments. (2015).

Eike Langbehn, Paul Lubos, and Frank Steinicke. 2018. Evaluation of locomotion techniques for room-scale vr: Joystick, teleportation, and redirected walking. In *Proceedings of the Virtual Reality International Conference-Laval Virtual*. ACM, 4.

Niels Christian Nilsson, Tabitha Peck, Gerd Bruder, Eri Hodgson, Stefania Serafin, Mary Whitton, Frank Steinicke, and Evan Suma Rosenberg. 2018. 15 Years of Research on Redirected Walking in Immersive Virtual Environments. *IEEE computer graphics and applications* 38, 2 (2018), 44–56.

Tabitha C Peck, Henry Fuchs, and Mary C Whitton. 2011. An evaluation of navigational ability comparing Redirected Free Exploration with Distractors to Walking-in-Place and joystick locomotion interfaces. In *2011 IEEE Virtual Reality Conference*. IEEE, 55–62.

Bernhard E Riecke and Jan M Wiener. 2007. Can people not tell left from right in VR? Point-to-origin studies revealed qualitative errors in visual path integration. In *2007 IEEE Virtual Reality Conference*. IEEE, 3–10.

Michael Lilienthal Robert Kennedy, Normal Lane, Kevin Berbaum. 1993. Simulator Sickness Questionnaire - An Enhanced Method for Quantifying Simulator Sickness. *The International Journal of Aviation Psychology* 3, 3 (1993), 203–220.

Roy A Ruddle and Simon Lessels. 2006. For efficient navigational search, humans require full physical movement, but not a rich visual scene. *Psychological Science* 17, 6 (2006), 460–465.

Roy A Ruddle and Simon Lessels. 2009. The benefits of using a walking interface to navigate virtual environments. *ACM Transactions on Computer-Human Interaction (TOCHI)* 16, 1 (2009), 5.

Thomas Schubert. 2002. Real and Illusory Interactions Enhance Presence. *Presence: Teleoperators & Virtual Environments* 11, 4 (2002), 425–434.

Salvar Sigurdarson, Andrew P Milne, Daniel Feureissen, and Bernhard E Riecke. 2012. Can physical motions prevent disorientation in naturalistic VR? In *2012 IEEE Virtual Reality Workshops (VRW)*. IEEE, 31–34.

Frank Steinicke, Gerd Bruder, Jason Jerald, Harald Frenz, and Markus Lappe. 2010. Estimation of detection thresholds for redirected walking techniques. *IEEE transactions on visualization and computer graphics* 16, 1 (2010), 17–27.

Evan Suma, Samantha Finkelstein, Myra Reid, Sabarish Babu, Amy Ulinski, and Larry F Hodges. 2009. Evaluation of the cognitive effects of travel technique in complex real and virtual environments. *IEEE Transactions on Visualization and Computer Graphics* 16, 4 (2009), 690–702.

Evan A Suma, Seth Clark, David Krum, Samantha Finkelstein, Mark Bolas, and Zachary Warte. 2011a. Leveraging change blindness for redirection in virtual environments. In *2011 IEEE Virtual Reality Conference*. IEEE, 159–166.

Evan A Suma, David M Krum, Samantha Finkelstein, and Mark Bolas. 2011b. Effects of redirection on spatial orientation in real and virtual environments. In *2011 IEEE Symposium on 3D User Interfaces (3DUI)*. IEEE, 35–38.

Evan A. Suma, Zachary Lipps, Samantha Finkelstein, David M. Krum, and Mark Bolas. 2012. Impossible spaces: Maximizing natural walking in virtual environments with self-overlapping architecture. 18, 4 (2012), 555–564.

Martin Usoh, Kevin Arthur, Mary C. Whitton, Rui Bastos, Anthony Steed, Mel Slater, and Frederick P. Brooks Jr. 1999. Walking> walking-in-place> flying, in virtual environments. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 359–364.

Bodo Winter. 2013. Linear models and linear mixed effects models in R with linguistic applications. *arXiv preprint arXiv:1308* (2013).

Catherine A Zanbaka, Benjamin C Lok, Sabarish V Babu, Amy Catherine Ulinski, and Larry F Hodges. 2005. Comparison of path visualizations and cognitive measures relative to travel technique in a virtual environment. *IEEE Transactions on Visualization and Computer Graphics* 11, 6 (2005), 694–705.