Mitigating Incorrect Perception of Distance in Virtual Reality through Personalized Rendering Manipulation

Alex Peer* Kevin Ponto†

Virtual Environments, Wisconsin Institute for Discovery
University of Wisconsin-Madison, USA

Abstract

Viewers of virtual reality appear to have an incorrect sense of space when performing blind directed-action tasks, such as blind walking or blind throwing. It has been shown that various manipulations can influence this incorrect sense of space, and that the degree of misperception varies by person. It follows that one could measure the degree of misperception an individual experiences and generate some manipulation to correct for it, though it is not clear that correct behavior in a specific blind directed action task leads to correct behavior in all tasks in general. In this work, we evaluate the effectiveness of correcting perceived distance in virtual reality by first measuring individual perceived distance through blind throwing, then manipulating sense of space using a vertex shader to make things appear more or less distant, to a degree personalized to the individual's perceived distance. Two variants of the manipulation are explored. The effects of these personalized manipulations are first evaluated when performing the same blind throwing task used to calibrate the manipulation. Then, in order to observe the effects of the manipulation on dissimilar tasks, participants perform two perceptual matching tasks which allow full visual feedback as objects, or the participants themselves, move through space.

Index Terms: I.3.7 [Computing Methodologies]: Graphics Utilities—Virtual reality

1 Introduction

Immersive display technologies that drive virtual, augmented, and mixed reality systems have been shown to induce distance misperceptions relative to real environments [13]. This misperceptions are often underestimations, earning this effect the name *distance compression* (DC). These misjudged distances have been cited as a concern for training applications in virtual reality (VR) [18], and have been shown to have some capability to transfer to behavior in real environments [1,2,9,12,16,17,19]. There is also some evidence that these misperceptions vary by viewer [7,10].

One method to "fix" the discrepancies between the perceived and intended virtual environments is to stretch the world out away from the participant, such that perceived distances match intended [7]. However, Peer and Ponto [7] found that while stretching the world by a fixed amount resulted in closer to real-world performance for a blind throwing task, it did not resolve distance compression to the degree expected. This work aims to further explore this result by changing the technique in two key ways. First, warps are calculated on an individual basis based on a blind throwing task that has been shown to accurately measure perceived distances [7, 8]. Following the suggestions of [8], an individual's distance compression is taken as the difference in task performance when viewing a real environment (RE) and a virtual environment (VE), in a measure called Relative Percent Error (RPE). Secondly, two types of warps are explored: one with a linear multiplier based on distance ($warp_a$), and a second which also includes a static offset ($warp_{ab}$). Conceptually, this offset is similar to the PO_z attribute referenced by Ponto et al. [10], in which the cameras representing the virtual eyes are placed at some distance inside of the participant's head.

However, it was noted by Peer and Ponto [7] that these warps may induce noticeable distortions. Though that work reported no noticeable adverse effects, it may be that these distortions aren't noticeable when performing blind throwing tasks, but would have an effect on other tasks; in general, the strategy of "fixing" distance compression using a single blind directed action task might "overfit" a manipulation to this single task, to the possible detriment of others. In this work, we observe the effect of the personalized warps on two perceptual matching tasks, similar to those used by Ponto et al [10].

This leads to the following hypotheses:

- **H1**: A personalized warp should reduce error in distances perceived when viewing virtual environments
- **H2**: The extra component of the model for the $warp_{ab}$ condition will improve performance over the $warp_a$ condition
- **H3**: Warps do not degrade performance in other perceptual tasks in the virtual environment

2 Related Work

Renner et al. [13] present a survey of the work on distance compression in VR, finding an average reported underestimation of 26% across 30 papers. As most papers do not compare real and virtual task performance, this average may overestimate the error [8].

Peer and Ponto [7] employed a technique they called "perceptual space warp", which attempted to mitigate distance compression by stretching the presented scene forward along the view axis by applying some multiplier to vertices in the z direction. They found a multiplier of 1.4 to induce only half of the expected change, but propose personalized multipliers might fare better.

Ponto et al. [10] attempted to use perceptual matching tasks to calibrate eye position on two axes to the individual participant, and establish a relationship between their eye-pose parameters, their perceptual matching tasks, and perceived distance. One eye pose parameter discussed in [10] may function similarly to the idea of changing the modeled center of the eye in VR rendering as discussed in [3].

Minification of the presented image through manipulating geometric field-of-view (gFOV) [4, 5, 15] has been shown to influence perceived distance. The warp multiplier used in this work has a similar visual effect to minification [7]. One issue with this approach is that field-of-view is a device specific feature, meaning that a calibration may need to be generated for both device and viewer.

The presented work attempts to mitigate perceptual issues by warping the space around the user, as first proposed Peer and Ponto [7]. In that work all users were assigned the same warp multiplier, meant to remove the effects of distance compression for

*e-mail: alex.peer@wisc.edu
†e-mail: kbponto@wisc.edu

a throwing task. While this warp multiplier was shown to significantly reduce distance compression, the effect was less pronounced than anticipated. The current work improves upon this previously presented method by:

1. Creating a per-person warp multiplier that attempts to match virtual world performance to real-world performance

2. Creating a warp that includes a multiplier as well as an offset

These methods are further tested against perceptual tests so as to examine the effect these multipliers have on a broader portion of users' experience.

3 METHOD

This experiment attempts to evaluate the distance misperception experienced by individuals in VR, and develop a customized set of parameters to drive a manipulation that elicits correct perceived distance. The effects of this attempted correction are then evaluated using three tasks.

3.1 Warps

We manipulate the scene by applying a vertex shader that shifts the position of things in the scene relative to the user's position and gaze direction. This is in the spirit of the perceptual space warping introduced in [7]; here, we'll simply call such manipulations *warps*.

Two different warps were tested in this study.

Warp$_a$: was determined by creating a linear regression based on the offset between the distances a participant threw to between the real and virtual conditions. The regression was forced to pass through the origin, meaning that it was assumed that objects at very small distances from the participant would not show distance effects. This warp is implemented using a vertex shader that transforms vertices by increasing their distance from the viewer on the viewspace z-axis. Formally, we define this transformation as:

$$V_{out} = M^{-1} \times ((M \times V_{in}) \times (1, 1, w_a, 1)) \qquad (1)$$

where V_{in} is an unwarped vertex position in 3-space and V_{out} is the warped resulting position, M is the modelview matrix, and w_a is the *warp multiplier*. By modifying w_a, one adjusts the magnitude and direction of the warp; a w_m of 1 causes no change, < 1 pulls objects closer, and > 1 pushes them further away.

Warp$_{ab}$: used a similar method as was used to determine $warp_a$, but added an additional component for world space offset. To accomplish this, the regression model was no longer forced to go through the origin, meaning objects at near distances could be offset via the warp. The warp is implemented using a vertex shader, which can formally be described as:

$$V_{out} = M^{-1} \times ((M \times V_{in}) \times (1, 1, w_a, 1) + (0, 0, w_b, 0)) \qquad (2)$$

where V_{in} is an unwarped vertex position in 3-space and V_{out} is the warped resulting position, M is the modelview matrix, and w_a is the *warp multiplier* and w_b is distance offset. As w_b is applied to all vertices regardless of their distance from the participant, one would hope that this value would be small; there is some question as to whether the position of a viewer's eyes is modeled correctly in current VR rendering [3, 10], and this might compensate for this sort of subtle misalignment.

The warps in this experiment are meant to match virtual to real-world performance. To this end, the regressions were fit to perceived distance in virtual environments, relative to mean perceived distance in real environments. To do this, perceived distance in the two environments was measured using a blind throwing task (see §3.2.1), yielding a perceived distance at given target distance, where target distances are modified by ± 0.5m from the three base target distances (2m, 3m, 4m), as described in §3.5. The real environment measurements were then binned at three intervals – the three base target distances – and the mean error in perceived distance for each RE base target distance was calculated. Finally, VE measurements were adjusted by the mean error of their matching RE base target distance; regressions were fit to this adjusted data to yield warp parameters with the intention of shifting perceived distance in the virtual environment to match that of the real environment.

3.2 Tasks

The experiment asked participants to perform three tasks: blind throwing, plank leveling (*planks*), and block squaring (*cube*). Blind throwing was used to measure the degree of error in perceived distance; all three were used to evaluate the effects of the warp manipulations.

3.2.1 Blind Throwing

We follow the protocol for blind throwing as established in works by Sahm et al [14] and Peer and Ponto [7, 8]; in [8] in particular, it was shown to be a directed action task where distances are perceived differently in virtual and real environments, and real world performance showed nearest to zero error. When performing this task, participants begin with their eyes closed. A sound then prompts them to open their eyes and view a triangular target placed at some distance between 1.5m and 4.5m, in either the real or virtual environment. After three seconds, participants are prompted to close their eyes, and throw a beanbag towards the target. The experimenter then registered the landing position of the target, removed the beanbag, and the next target was prepared – manually placed in RE conditions, using an LED light strip positioned to one side of the throwing space to indicate both the next target position and the position of a controller held near the target, such that matching the two lights indicated an aligned target.

The primary error measure used for blind throwing is relative percent error (*RPE*) as described in [8]. First, perceived distance is taken as the distance of a beanbag's landing point, projected onto the axis in the ground plane running forward in the direction the participant has been asked to face towards – the axis on which targets are placed. Error in perceived distance is taken to be the difference between perceived and target distance, such that error is negative when perceived distance is less than target distance, and percent error is this value divided by the target distance; for RPE, we adjust the target distance by the mean error seen in real environment conditions. For this experiment, this RE mean error is binned by the three base target distances described in §3.5; VE measures are adjusted by the RE mean error belonging to the base target distance nearest to the VE target distance.

3.2.2 Planks

The planks task mirrors the perceptual matching task used for distance estimation in Ponto et al. [10], in which perceived slope is used as a proxy measure for distance perception [6, 11]. This task is intended to isolate stereo vision as a depth cue. In this style of tasks, participants are asked to level a tilted plane.

As in [10], 5 planks with a depth of 0.52m were placed at one meter intervals, with a random perturbation of $\pm 0.25m$. Planks are positioned 0.72m above the floor. Our implementation differs in that the planks have no height (0.001m) and infinite width (10000m), in hopes of further isolating the desired depth cue.

The planks are "tilted" by some angle, such that an angle of 0 aligns them along a flat plane, a negative angle orients them as a staircase leading down, and a positive angle as a staircase leading up. The staircase effect is due to forcing the planks to remain parallel to the desired flat plane, as [10] found that aligning the planks to the tilted plane indicated by the selected angle induced unintended

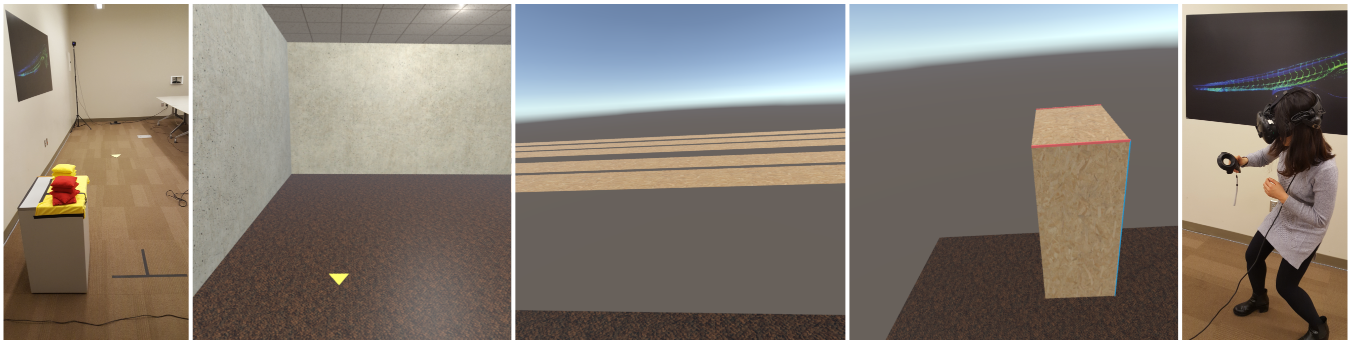

Figure 1: The environments and tasks used during the experiment. From left to right: The real environment for throwing; the virtual throwing environment; the planks task; the cube task; a collaborator demonstrating the cube task.

depth cues through motion. In this work, the intended z position of individual planks is preserved by choosing their y position by casting a ray from their centers in the y direction, and positioning them at the intersection of this ray and the tilted plane.

When executing this measure, participants stood in place and held a tracked controller. When holding the trigger, moving the controller up or down would work as a sort of open-air joystick, causing the planks orientation to change by some constant increment relative to the distance of the controller on the y-axis, relative to its position when the trigger started being held down. This was described to participants as holding the trigger and "pulling" the planks up or down. The intent is to not provide too a direct mapping between the planks' movement and the participant's proprioceptive sense of distance.

At the beginning of each planks trial, the planks were tilted by a random angle between -5 and 5 degrees. The error measure for the planks task is the angle selected; any deviation from 0 is error. For the sake of observing whether a warp makes a task more difficult, we are concerned less with the directionality of the error as with its magnitude, so we also investigate absolute error measured as the absolute value of degrees from 0.

3.2.3 Cube

This task mirrors the perceptual matching task used to judge perception of shape by Ponto et al [10]. Participants are presented with a floating box with each dimensions randomly sized, between 0.15m and 0.6m. One dimension is held fixed, and participants are asked to manipulate the other two such that all sides are the same size; that is, they are asked to make the box into a cube. Ponto et al. [10] found that this task can be affected by manipulating the way a virtual environment is presented.

In our implementation, the box floated 1m above the ground. The z dimension was held fixed. Colored guides were placed on the two top edges running along the x dimension (red), and two opposing edges running along the y edges (blue), such that one edge of each color should always be in view as participants moved through the space. A similar open-air joystick control scheme as used for the planks task was used here, with movement along the box's x and y axes causing the cube's respective dimension to grow or shrink; participants were told to "pull" in the red or blue direction to make the cube grow or shrink in that direction.

The error metric for the cube task is as in [10]: an L2 norm, calculated thusly:

$$Error = \sqrt{(l_x - l_z)^2 + (l_y - l_z)^2} \tag{3}$$

Where l_x, l_y, l_z is the cube's length along the subscripted x, y, and z axes, respectively.

3.3 Participants

13 participants were recruited from a local university campus. One participant was removed from the current analysis, as extreme responses to the planks task and experimenter observation during execution suggest they may have experienced difficulty in completing the task. The remaining 12 participants ranged in age from 20 to 28 (M: 23, SD: 2.5), 6 female and 6 male.

3.4 Materials

The experiment was administered in an 8m by 6m conference room, with a roughly 8m by 3m lane cleared for the experiment. The lighthouse tracking system of an HTC Vive was set to track a 6m long portion of this space, with the lighthouses positioned roughly 6.6m apart, connected by an optical sync cable and set to optical sync master and slave modes. Vive controllers were used as input devices during the two perceptual matching tasks, and to register beanbag landing positions during throwing tasks. An LED lightstrip was positioned to the left side of the space, and was used to indicate to the experimenter where real-world throwing targets should be placed, and whether Vive controllers were tracking accurately during beanbag landing registration. The LEDs were positioned roughly 2.5cm apart.

3.5 Design

This experiment followed a within-participant design, with each participant exposed to all conditions.

Three rounds of blind throwing were performed. Conditions of two factors were presented: *Viewing Condition* and *Target Distance*. *Viewing Condition* has four levels: the real environment (*real*), a virtual environment with no warp applied (*no warp*), a VE with a $warp_a$applied, and a VE with a $warp_{ab}$applied. *Target Distance* has three levels: 2, 3, and 4 meters. All target distances were used in all rounds of throwing. Target distances were always presented in five randomly ordered sets of the three levels, such that each target distance was presented five times. The actual distance presented was adjusted slightly from the base target distance, by $\pm$ 0.5m.

The first two rounds of throwing were pretests, *real* then *no warp*, over all distances for a total of 15 trials each. The third round of throwing was a post-test, and randomly presented both $warp_a$ and $warp_{ab}$ conditions over all distances, for a total of 30 trials, or 15 for each *Viewing Condition*.

Two perceptual matching tasks were performed. Both presented conditions of one factor, *Warp Method*, with three levels: *none*, *a*, *ab*. Both tasks presented scenes using the three methods in random order. Both tasks presented each warp method 5 times, for a total of 15 trials each.

3.6 Procedure

After reviewing and signing the consent form, participants were shown to the testing area. They were then asked to stand with their heels on a line on the ground made of tape, marking the zero distance point for the throwing task. Participants were asked to face forward along the axis on which targets would be placed, marked by another line of tape running forward between their feet. These tape lines were reproduced in the virtual environment used for throwing.

The throwing protocol was explained and the sounds indicating when to open and close their eyes were demonstrated. A target was placed near the center of the space, and participants were allowed to practice throwing beanbags at this target until they indicated confidence in their ability.

Participants were then asked to close their eyes, and performed the real environment round of pretest blind throwing. After this, the experimenter explained the use of the HMD, providing instructions on adjusting its fit. Once the headset was adjusted for comfort and clarity, the virtual environment pretest throwing was performed.

Once both pretest phases were complete, the per-participant parameters for the corrective manipulations were calculated. The post-test VE throwing phase was then executed.

Next, participants were asked to remove the HMD, and the use of the controller and the goal of the planks task was explained. Participants then wore the headset and performed the planks phase. After this, the cube phase was explained, then performed.

After completing the main experiment participants filled out a short demographic survey, and were asked to complete a colorblindness test using Ishihara plates, and a stereoblindness test using a random-dot stereogram. Finally, a post-experiment interview and debrief was then performed.

4 Results

4.1 Analysis

4.1.1 Throwing Task

Viewing Condition	N	Percent Error (PE) Mean	SD	Relative Percent Error (RPE) Mean	SD
real	180	-3.77	9.20	–	–
no warp	180	-10.81	12.49	-6.96	13.71
warp_{a}	180	-6.59	10.84	-2.79	10.74
warp_{ab}	180	6.23	20.23	10.03	19.83

Table 1: Summary of throwing trials.

Effect	On	Df_n	Df_d	F	p	Sig.	η_P^2
Condition (C)	PE	3	33	8.87	.008	**	.45
Distance (D)	PE	2	22	30.47	$<.001$	***	.73
C *x* D	PE	6	66	8.11	.001	**	.42
Condition (C)	RPE	2	22	10.55	.006	**	.49
Distance (D)	RPE	2	22	20.54	$<.001$	***	.65
C *x* D	RPE	4	44	2.86	.035	*	.21

Table 2: ANOVA of throwing trials. Df_n and Df_d are the numerator and denominator degrees of freedom, F is the F-value, p is conditional probability of the F-test, η_P^2 is partial eta-squared. PE is percent error, RPE is percent error in virtual environment trials relative to real environment trials.

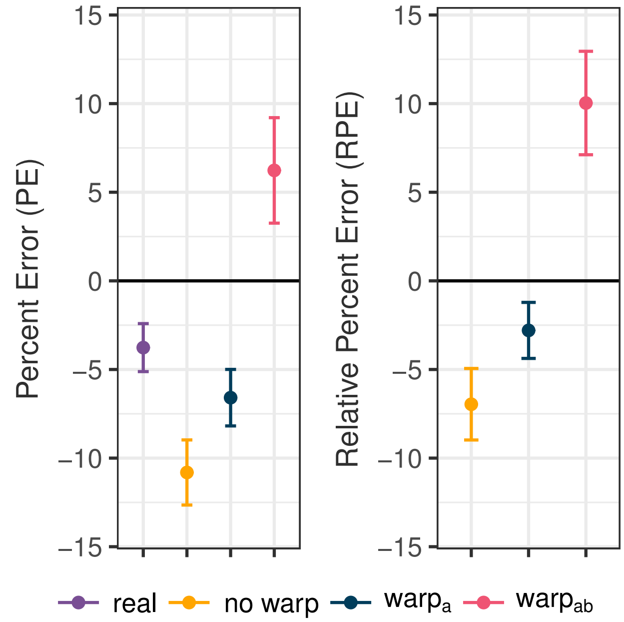

Figure 2: For the throwing task in each viewing condition, mean percent error (left) and mean relative percent error (right). Error bars represent 95% confidence intervals. Negative error represents underestimation, positive error represents overestimation.

Conditions		Metric	Df	t	p	Sig.
real	no warp	PE	179	-8.98	$<.001$	***
no warp	warp_{a}	RPE	179	-4.4355	.005	**
no warp	warp_{ab}	RPE	179	-10.432	$<.001$	***
warp_{a}	warp_{ab}	RPE	179	-10.691	$<.001$	***

Table 3: Pairwise t-tests comparing conditions from throwing trials. Holms-Bonferroni was used to correct for multiple comparisons.

Results of repeated-measures ANOVAs on the effect of viewing condition and target distance on percent error and relative percent error can be seen in Table 2; all factors show significant differences between groups. The effects of target distance are somewhat surprising and will be revisited in the discussion, but for the sake of space will not be further examined here.

Pairwise t-tests on viewing condition, as shown in Table 3, show that the RE and VE pretest conditions (*real* and *no warp*) are significantly different, which suggest participants exhibited distance compression. Each combination of *no warp*, $warp_a$, and $warp_{ab}$ were also found to be significantly different from one another, suggesting warp had some effect on perceived distance.

Table 1 and Figure 2 summarize the throwing trials. In *no warp* conditions the mean percent error (-10%) and relative percent error (-7%) are lower than seen elsewhere – percent error of 25% [13], or relative percent error of -18% [8]. This suggests participants experienced lower than average distance compression. Changes in relative percent error between *no warp* to $warp_a$ conditions show mean error roughly halving, indicating the $warp_a$ condition had a positive effect on perceived distance; this shows support for **H1** in $warp_a$ conditions. The $warp_{ab}$ conditions show significant overestimation, which suggests they did not have the intended effect; this does not support **H2**, nor **H1** for $warp_{ab}$ conditions.

4.1.2 Perceptual Matching Tasks

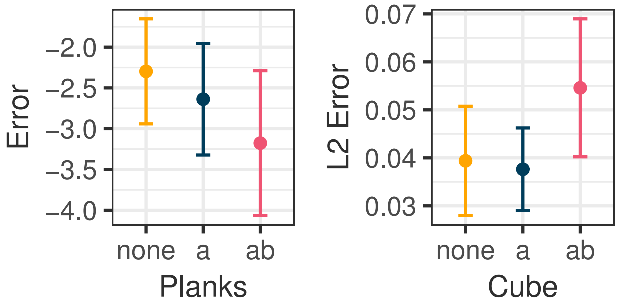

Figure 3: Mean error when performing the two perceptual tasks. On the left, degrees error when performing the planks task. On the right, L2 error when performing the cube task. Error bars represent 95% confidence intervals.

		Planks (degrees)		Cube (m)	
		Error		L^2 Error	
Warp	N	Mean	SD	Mean	SD
none	60	-2.30	2.49	0.039	0.044
a	60	-2.64	2.64	0.038	0.033
ab	60	-3.18	3.43	0.055	0.056

Table 4: Summary of the perceptual matching tasks' results.

Task	Metric	Df_n	Df_d	F	p	Sig.	η_P^2
Planks	Error	2	22	1.38	.269		.11
Cube	L2	2	22	3.98	.033	*	.27

Table 5: ANOVA results for the two perceptual matching tasks. Headings as in Table 2, with the exception of dependent variable here being the error metric from the referenced task.

		Task	Df	t	p	Sig.
none	a	Cube	59	0.28	.778	
none	ab	Cube	59	-1.88	.064	
a	ab	Cube	59	-2.26	.027	*

Table 6: Pairwise t-tests comparing warp methods during the cube task. Holms-Bonferroni was used to correct for multiple comparisons.

Results of repeated-measures ANOVAs on the effect of warp method on the error metrics of the perceptual matching tasks can be seen in Table 5. Warp method is found to have a significant effect on the cube task's L2 error, but not on the planks task's degree error; this lack of significance lends some support to **H3**.

Pairwise t-tests comparing the error metrics of the various warp methods for the cube task are shown in Table 6; only $warp_a$ and $warp_{ab}$ conditions are found to be significantly different; this further supports **H3**.

A summary of the observed error in perceptual matching tasks can be found in Table 4, as well as Figure 3. The planks task sees a trend towards increasing error and variance in both $warp_a$ and $warp_{ab}$ cases, relative to the *no warp* case; this slightly suggests **H3** may not hold for the planks task. All planks conditions show significantly larger errors than seen in Ponto et al [10]. The cube task sees similar amounts of error in *no warp* and $warp_a$ cases; this supports **H3** in the $warp_a$ conditions. $Warp_{ab}$ shows an increase of 0.016m mean error, as well as increased variance; this suggests **H3** may not hold in $Warp_{ab}$ conditions. The cube task shows mean errors of a similar magnitude as Ponto et al [10].

4.2 Discussion

In regards to the hypotheses proposed in §1, the $warp_a$ technique shows promising support for **H1** and **H3**; $warp_a$ reduced distance compression in the throwing task and showed no significant influence in the two perceptual tasks. The $warp_{ab}$ technique, however, not only increased error in throwing, it caused overestimation – neither **H1** nor **H2** seem to hold for this technique; oddly, though $warp_{ab}$ does increase mean error in the perceptual tasks, it had no statistically significant influence, suggesting **H3** may still hold for when employing $warp_{ab}$. Overall, $warp_a$ may be a viable technique for mitigating distance misperception in VR; $warp_{ab}$, as implemented here, is not, but allows observation of the effect of extreme warp on the three tasks.

		$Warp_a$		$Warp_{ab}$	
parameter	N	Mean	SD	Mean	SD
a	12	1.107	0.069	1.281	0.209
b	12	-	-	-0.343	0.276

Table 7: Summary of warp parameters chosen to correct for each participant's differing amount of observed distance compression.

That $warp_{ab}$ had such a large detrimental effect on throwing trials was unexpected. To fit as a proxy of eye position as in the model proposed by Ponto et al. [10], we would hope that the values of the offset from $warp_{ab}$ be small, on the order of centimeters in magnitude. However, as can be seen in Table 7, calibration from the linear regression selected offsets which were in general quite large, in tens of centimeters; the mean w_b of -0.34 would shift the user's position backwards by 34cm. Not only is this well out of the range of candidate eye positions, it would result in significant changes to the scene at near distances when the *a* multiplier does less to cancel out the effect. For example, when looking directly at one's feet, the ground would appear to be almost 34cm higher.

This likely created a situation in which closer objects were overly displaced, as shown in the 2m column of Figure 4. It is interesting to note that while fairly extreme overestimation was shown in the throwing task, this did not translate to selection of significantly different negative angles of inclination in the planks task as would be indicated by Ponto et al [10]. This may suggest that the throwing task and planks task draw from different depth cues, or that the warp multiplier changes the cues used during one but not the other. This may also be due in part to differences between this experiment and that of Ponto et al. The earlier experiment used a CAVE, allowing for a visible physical reference object; the HMD used here did not. It may also be due to the differences in task implementations: the planks here are infinitely wide, rather than fixed width and jittered on the x-axis; the planks here have no apparent height, rather than 0.46m.

Several participants were observed looking along the planks to where the infinite length met the horizon, particularly under extreme $warp_{ab}$ conditions. Participants did not report this helping their judgments, however; when interviewed, it seemed that looking at an infinitely long stimulus that would not be influenced by the warp may have afforded a more comfortable or less cognitively demanding stimulus, though they hesitated to assign it to either when

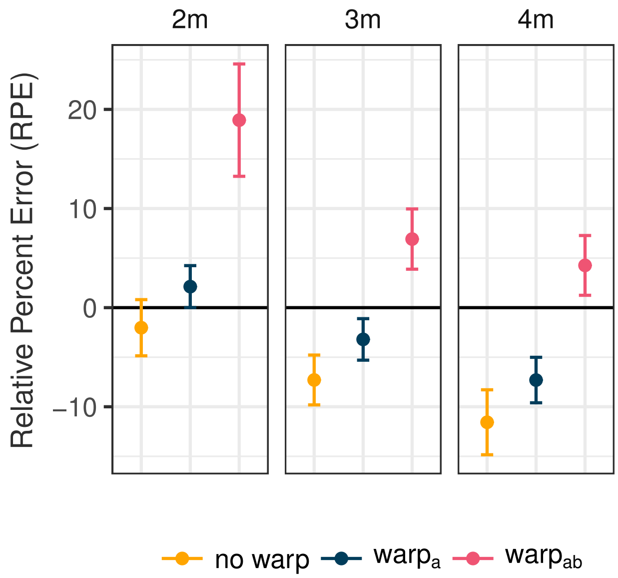

Figure 4: Mean relative percent error when performing the throwing task, split horizontally by the three base target distances. Error bars represent 95% confidence intervals.

asked directly. This gives some slight indication that extreme warp may have some detrimental effect outside of task performance as measured in this work, one that participants may not be completely aware of themselves but do seek to alleviate.

It is also interesting that despite $warp_{ab}$ inducing particularly strong overestimation at the closest distance of 2m, it did not have an effect of similar magnitude on the cube task, which was viewed from distances exclusively closer than 2 meters. This may suggest that under extreme warps, a different strategy was employed – either different depth cues were used, or cues entirely separate from depth were somehow employed. This suggests users might adapt, even to extreme warps. Participants did not report consistently using any alternative strategies. Also worth noting is that participants had no trouble walking around the cube in either warp condition.

While $warp_a$ showed throwing task performance that better matched real-world conditions, it does not induce fully matched performance. Looking again to Figure 4, it seems $warp_a$ performed better at shorter distances; this may indicate that the multipliers were biased by there being so little distance compression evident at the 2m distances in *no warp*, which is itself an anomaly. We see amounts of distance compression similar to Interrante et al. [1], who made the serendipitous discovery that a virtual environment matched to the real environment occupied by a participant appeared to induce greatly reduced underestimation – percent error on the order of -10%, in both real and virtual trials. The virtual environment used here was a spatially matched but sparse environment, in the style of Peer and Ponto [8], which work uses relative percent error and saw somewhat more underestimation (-18%) than we do here (-7%). It may be that the virtual environment used in the current work is somehow too good a match and runs afoul of the matched environment effect seen by Interrante et al., at least partially, and particularly at near distances. However, we do see more underestimation in virtual trials than real trials, and so our participants do seem to have experienced distance compression.

$Warp_a$ also caused no significant detriment in performance for either perceptual matching task; for the cube task, this suggests that $warp_a$ did not introduce distortions that interfered with the task. For the planks task, improved distance estimates in the throwing tasks would be expected to correspond to reduced error in planks trials; however, no significant difference was found, and the mean error is slightly worse in $warp_a$ conditions than *no warp*. Similarly to the $warp_{ab}$ results, this may be due to differences between our implementation of the task and those of previous works [10]. It may also indicate that the depth cues used by the planks task either are not or are negatively influenced by the warp manipulation; further work is needed to decide.

At 4m, the two warp methods seem to bracket a relative percent error of 0; further, $warp_a$ provides the least mean error at 2m, and steadily more as distance increases; $warp_{ab}$ provides less mean error that $warp_a$ at 4m, with $warp_{ab}$ error increasing as distance decreases. This may indicate that some combination of the two might be better than either alone, that some nonlinear or piecewise linear function might be optimal. Further work might also explore $warp_{ab}$ at longer distances to see if the trend of decreasing mean error continues; conversely, adding pre-test measures at even closer distances might prevent extreme w_b values.

5 Conclusion

Of the two warp methods tested, $warp_a$ seems the most immediately viable, as it improved distance estimates in throwing trials and did not disrupt estimates in the cube task. However, it did not significantly influence the planks task, which was expected to be influenced by the sense of depth manipulated by the warp. The $warp_{ab}$ method used resulted in extreme warp parameters, which caused overestimation in throwing trials, but had no significant influence on either perceptual task. Overall, the basic premise of mitigating distance misestimation in virtual reality using personalized manipulations does seem viable. More work is needed to explore means of calibrating and implementing warps across a range of distances, and to observe the effects of warps on other classes of tasks than those explored here. Manipulations other than warps may also be possible, with their own strengths and caveats.

References

[1] V. Interrante, B. Ries, and L. Anderson. Distance perception in immersive virtual environments, revisited. In *IEEE Virtual Reality Conference (VR 2006)*, pages 3–10. IEEE, 2006.

[2] V. Interrante, B. Ries, J. Lindquist, M. Kaeding, and L. Anderson. Elucidating factors that can facilitate veridical spatial perception in immersive virtual environments. *Presence: Teleoperators and Virtual Environments*, 17(2):176–198, 2008.

[3] J. A. Jones, D. Edewaard, R. A. Tyrrell, and L. F. Hodges. A schematic eye for virtual environments. In *2016 IEEE Symposium on 3D User Interfaces (3DUI)*, pages 221–230. IEEE, 2016.

[4] S. A. Kuhl, W. B. Thompson, and S. H. Creem-Regehr. Hmd calibration and its effects on distance judgments. *ACM Transactions on Applied Perception (TAP)*, 6(3):19, 2009.

[5] B. Li, R. Zhang, A. Nordman, and S. A. Kuhl. The effects of minification and display field of view on distance judgments in real and hmd-based environments. In *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception*, pages 55–58. ACM, 2015.

[6] T. L. Ooi, B. Wu, and Z. J. He. Distance determined by the angular declination below the horizon. *Nature*, 414(6860):197, 2001.

[7] A. Peer and K. Ponto. Perceptual space warping: Preliminary exploration. In *IEEE Virtual Reality (VR)*, pages 261–262, March 2016.

[8] A. Peer and K. Ponto. Evaluating perceived distance measures in room-scale spaces using consumer-grade head mounted displays. *IEEE 3D User Interfaces (3DUI)*, 2017.

[9] J. M. Plumert, J. K. Kearney, J. F. Cremer, and K. Recker. Distance perception in real and virtual environments. *ACM Transactions on Applied Perception (TAP)*, 2(3):216–233, 2005.

[10] K. Ponto, M. Gleicher, R. G. Radwin, and H. J. Shin. Perceptual calibration for immersive display environments. *Visualization and Computer Graphics, IEEE Transactions on*, 19(4):691–700, 2013.

[11] D. R. Proffitt, M. Bhalla, R. Gossweiler, and J. Midgett. Perceiving geographical slant. *Psychonomic bulletin & review*, 2(4):409–428, 1995.

[12] D. R. Proffitt, J. Stefanucci, T. Banton, and W. Epstein. The role of effort in perceiving distance. *Psychological Science*, 14(2):106–112, 2003.

[13] R. S. Renner, B. M. Velichkovsky, and J. R. Helmert. The perception of egocentric distances in virtual environments-a review. *ACM Computing Surveys (CSUR)*, 46(2):23, 2013.

[14] C. S. Sahm, S. H. Creem-Regehr, W. B. Thompson, and P. Willemsen. Throwing versus walking as indicators of distance perception in similar real and virtual environments. *ACM Transactions on Applied Perception (TAP)*, 2(1):35–45, Jan. 2005.

[15] F. Steinicke, G. Bruder, S. Kuhl, P. Willemsen, M. Lappe, and K. Hinrichs. Natural perspective projections for head-mounted displays. *IEEE transactions on visualization and computer graphics*, 17(7):888–899, 2011.

[16] D. Waller and A. R. Richardson. Correcting distance estimates by interacting with immersive virtual environments: Effects of task and available sensory information. *Journal of Experimental Psychology: Applied*, 14(1):61, 2008.

[17] B. G. Witmer and W. J. Sadowski. Nonvisually guided locomotion to a previously viewed target in real and virtual environments. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 40(3):478–488, 1998.

[18] R. H. Wright. Virtual reality psychophysics: Forward and lateral distance, height, and speed perceptions with a wide-angle helmet display. Technical report, DTIC Document, 1995.

[19] C. J. Ziemer, J. M. Plumert, J. F. Cremer, and J. K. Kearney. Estimating distance in real and virtual environments: Does order make a difference? *Attention, Perception, & Psychophysics*, 71(5):1095–1106, 2009.