Perceptual Space Warping: Preliminary Exploration

Alex Peer* Kevin Ponto†

Living Environments Lab
University of Wisconsin-Madison

ABSTRACT

Distance has been shown to be incorrectly estimated in virtual environments relative to the same estimation tasks in a real environment. This work describes a preliminary exploration of *Perceptual Space Warping*, which influences perceived distance in virtual environments by using a vertex shader to warp geometry. Empirical tests demonstrate significant effects, but of smaller magnitude than expected.

Index Terms: H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities ; I.3.7 [Computing Methodologies]: Graphics Utilities—Virtual reality

1 INTRODUCTION

Distance compression (DC), the perceived underestimation of ego-centric distances, is a well researched phenomenon in *virtual environments* (VEs). Renner et al. found an average reported underestimation of 26% across 30 papers[2]. While many researchers have noted the effects of distance compression, and many possible influences have been explored, the principle factors responsible for this phenomenon remain largely unknown.

This paper presents *Perceptual Space Warping* (PSW), which aims to mitigate distance compression by altering the VE in a shader during runtime. PSW results in a final image similar to the geometric minification shown by [1]; however, PSW can be used without knowledge of device screen parameters or calibrated FOV, and at a different place in the rendering pipeline. Although not explored here, PSW could allow for piecewise linear or nonlinear warps; it also influences depth directly, allowing it to work in concert with other techniques that would not be affected by image-space manipulations, such as depth of field. A potential disadvantage: at larger warp modifiers, PSW introduces a distortion we call "swim", which causes the vague sensation, but not the actual visuals, of surfaces rippling or bending. Also, at very large warp modifiers, the scene is noticeably stretched.

2 METHOD

PSW 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:

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

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_m is the *warp multiplier*. By modifying w_m, 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. This shader was implemented in the Unity game engine. The Oculus Rift DK2 head-mounted display (HMD) was used to display the VE, using the Oculus v0.6 SDK and Unity integration.

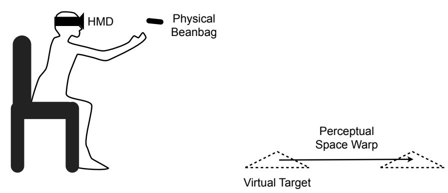

Figure 1: PSW moves objects relative to the viewer.

A blind throwing task was used to measure participant's perception of distance, similar to Sahm et al's [3]; we differ in that our participants are seated, and we use no blindfolds in Experiment 1.

3 EXPERIMENT 1

The goal of Experiment 1 (E1) is to confirm that our task does elicit and measure distance compression, and to explore the effect of PSW. This yields five key hypotheses: ***H1***: Participants will be able to throw bean bags accurately, in the favorable conditions of *real environment* (RE) with visual feedback. ***H2***: Participants will experience distance compression when viewing our VEs. ***H3***: PSW will have a significant and positive effect mitigating distance compression. ***H4***: The virtual reality environment will change the degree of effect of PSW. ***H5***: Target distance will not influence percent error of thrown distance to target.

10 participants were recruited during an outreach event held between our lab and a local nonprofit; due to the informal nature of the event, no demographics were gathered.

3.1 Design and Procedure

Participants proceeded through three phases in the following order: ***P0*** Practice Throwing, ***P1*** RE Throwing, ***P2*** VE Throwing.

In *P2*, 3 factors were investigated in the VE: ***warp***, the presence or absence of PSW (within, 2 levels: no warp, w_m of 1.4); ***environment***, the surrounding environment in which the target is presented (between, 2 levels: sparse, rich); ***distance***, the distance of the target from the participant (within, 11 levels: 0.1m intervals between 4m and 3m).

In *P2* participants threw once at each of the 11 distances in each warp condition, presented in random order, for a total of 22 throws in the VE. Each participant saw one environment, half assigned to sparse and half to rich. *P1* consisted of 11 real-world throws at distances of 3m (4 throws), 3.5m (3 throws), and 4m (4 throws), in random order. *P0* asked participants to practice throwing at a 3.7m target until they felt comfortable with the task. To maintain the brisk pace necessitated by the venue, no blindfolds were used in the real-world, and the screen was not blanked in the virtual.

3.2 Analysis

A 2(*warp*)x2(*environment*)x11(*distance*) ANOVA shows a significant increase in accuracy in warp 1.4 conditions ($F(1,8) = 69.7780$, $p < 0.0001$). No other significant main effect or interaction was found. This supports *H3* and *H5*, but not *H4*.

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

Participants threw accurately in the RE, (percent error: M=-3.6, SD = 8.1), and underthrew in VE no warp conditions (percent error: M=-32.3, SD=12.6), supporting *H1* and *H2*; t(9)=7.25, p=<0.0001.

Participants underthrew more than expected in VE warp 1.4x conditions (percent error: M=-18.1, SD=15.1; DC corrected: M=14.1, SD=5.1), an interesting result for *H3*.

3.3 Discussion

Support for *H1*, *H2*, and *H5* suggest our methods are viable for further experiments. Lack of support for *H4* suggests PSW is not reliant on (or influenced by) depth cues present in rich scenes. This is unexpected and may bear further investigation. Support for *H3* in general was not unexpected, but it is surprising that a 1.4 warp multiplier results in a significantly smaller mean correction of 14%, rather than the expected 40%.

4 Experiment 2

Experiment 2 (E2) investigates three potential confounds or enhancements suggested by E1: Swim distortion, Calibration, and w_m to Distance Mapping Exploration. E2 progressed as follows: ***P0*** practice throwing, ***P1*** throwing over intervals, ***P2*** calibration, ***P3*** calibrated throwing, ***P4*** real-world throwing and post-experiment interview. The practice phase *P0* was as in E1. *P4* saw real-world throwing to a target at 4m using the HMD as a blindfold; participants who were not accurate in *P4* were removed from the current analysis (2 removed). Participants were also screened for stereo sensitivity via a random dot stereogram test (3 removed). 9 participants total provide the data for our final analysis, recruited from the local campus (ages 19-35; 4 female). The throwing task in *P1* and *P3* was similar to that of E1, but the scene was hidden during the throw to better match [3]. Scenes were displayed for 5 seconds before the screen went blank.

Swim: As discussed in the introduction, PSW introduces a distortion we call *swim* which may change perception of depth with change of head pose. We created a second pose-independent warp shader that creates a warp valid for only a single head rotation. This eliminates the swim distortion for the controlled scenes in our experiment, but not the general case. Both *P1* and *P3* involve swim distortion as a factor (swim) of two levels: full swim, no swim.

Calibration: In *P2* we ask participants change w_m via joystick to visually align a virtual and real target, by repeatedly removing and replacing the HMD to view both VE and RE. The average of 4 trials is the participant's calibrated VM, which is used for blind throwing in *P3*.

Mapping: In *E1*, a w_m of 1.4 did not map to the distance expected. *P1* asks participants to throw at many w_ms, to better describe the relationship between w_m and perceived distance.

This yields the following specific hypotheses:

H1. Eliminating the swim distortion will improve blind throwing accuracy. ***H2***. Users can perform the calibration task of matching a virtual and real target. ***H3***. Using a user-calibrated warp multiplier will improve blind throwing accuracy. ***H4***. The optimal warp multiplier will be dependent on the individual participant. ***H5***. The relationship between warp multiplier and distance thrown is linear and monotonically increasing.

4.1 Design and Analysis

P1: 8 x 2 factorial, within-participant: 8 w_m (0.6 to 2.0 at intervals of 0.2) x 2 *swim* (full swim, no swim), 3 throws per condition. Target distance also varied; each participant threw to the same set of 48 distances between 3m to 4m (0.1m intervals). 48 throws total, random order, double blind.

P3. 2 x 11 factorial, within-participant: 2 *warp* (no warp, personal warp) x 2 *swim* (full swim, no swim) x 11 *distances* (3m to 4m at 0.1m intervals). 22 throws total, random order, double blind.

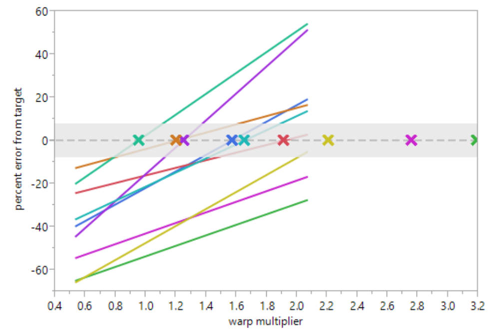

Figure 2: Linear fits of participants' throwing error over different w_ms. Xs indicate w_ms expected to yield 0 error for that participant.

P1: An 8 (w_m) x 2 (swim) repeated measures ANOVA showed only a significant main effect of warp multiplier (F(7,56)=37.95, p < 0.0001); no further significant main effect or interaction. This shows no support for *H1*. As shown in Figure 2, the shape of the graphs support *H5*, and the individual variance of the x-axis crossings where error is 0 supports *H4*.

P2: 7 of 9 calibrated w_ms are near or less than 1 (M=-0.03, SD=0.02; 4 <0.9, 4 ≈1, 2 >1.1). This does not support *H2*.

P3: A 2 (swim) x 11 (distance) repeated measures ANOVA showed no significant main effects or interactions. This shows no support for *H1*, though it further reinforces E1 *H5*. Calibrated w_ms result in significant percent throwing error (M=-28.9, SD=16.9), and little per-participant correction (M=-3.7, SD=6.2). This does not support *H3*.

4.2 Discussion and Conclusion

Swim: We see no evidence that removing swim influenced perception of distance. In post-experiment interviews, only two participants mentioned noticing swim when it was described to them. The throwing task used may not elicit enough head movement to fully show the effects of swim.

Calibration: In P2, only 2 participants selected w_ms that would mitigate DC (>1), and calibrated w_ms did little to correct DC in P3. Had participants chosen consistent w_ms, it might indicate they were selecting the correct calibration for the HMD; had we seen increased DC correction, we might conclude each participant perceived a unique miscalibration. In the absence of either, we suspect our calibration task was too difficult to perform accurately.

Mapping: The mapping seems to be fairly linear and monotonically increasing. Per-user differences in estimated w_m for 0 error suggests a per-user w_m may be worth exploring, though some are quite extreme (> 2) and may induce distortion.

Acknowledgements

This work was supported in part by the University of Wisconsin - Madison Fall Research Competition.

References

[1] 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.

[2] 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.

[3] 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, 2005.