

Self Localization in Virtual Environments using Visual Angles *

Scott A. Kuhl[†]
Augsburg College
Minneapolis, MN 55454

Karen T. Sutherland[‡]
Augsburg College
Minneapolis, MN 55454

Abstract

This work analyzes judgments of self localization in immersive virtual environments. A sparse environment was created and a set of localization tasks requiring that angle estimates be made to three landmarks was presented to two sets of subjects, one given feedback and one not. Results were compared to experiments in robot navigation as well as real world perceptual experiments. Feedback improved performance as did the configuration of the landmarks used for a task.

CR Categories: H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality

Keywords: Immersive virtual environment, localization, perception

1 Introduction

How spatial judgments in a virtual environment correspond to those in the real world remains an open question. A well understood correspondence would allow researchers in human perception to make use of results from experiments in the more easily controlled virtual environment. Understanding the nature of the correspondence would also help in rendering more realistic virtual environments and graphical simulations. This work focuses on self localizing in immersive virtual environments using angular measurements to three landmarks. The goals of the research were to learn more about how people make spatial judgments in virtual environments and compare the results of those judgments with similar work done with human subjects in the real world as well as geometric analysis that has been done in robot navigation.

A thorough analysis of error in robot self localization using visual angle measurements to three landmarks in the environment [Sutherland 1993; Sutherland 1994] has shown that, for any given error in visual angle measure, landmark configuration affects error in localization. It is well known that accurate measure of the angles subtended to three point landmarks in the environment is sufficient

*This work was supported by National Science Foundation grant #0080986 (ITR)

[†]e-mail: scott@kuhlweb.com

[‡]e-mail: suther@navigation.augsburg.edu

Copyright © 2004 by the Association for Computing Machinery, Inc. 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 commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1 (212) 869-0481 or e-mail permissions@acm.org.

© 2004 ACM 1-58113-884-9/04/0006 \$5.00

to determine self location on a map on which those landmarks have been identified. The one exception is if the landmarks and the observer's location all lie on a single circle, in which case, there is not a unique solution. When error in measurement is introduced, the error in localization is continuous, becoming more severe as the configuration approaches single circle. Thus, configurations which produce a large change in angle measure for a small movement of the observation point are considered better for a robot navigating with visual sensors. Furthermore, if a bound on the error in measurement is known, the area in which the viewpoint will lie ("area of uncertainty") can be accurately determined. The best configurations are those in which the center landmark is closer to the viewpoint than the outer two. Configurations can be compared using a measure of "goodness" which ranks them according to the localization they produce for any amount of sensor error.

Wagner et al [Wagner et al. 1994] ran a series of experiments based on the above results with subjects determining their own location on a map using visual angle measures to three landmarks in an outdoor environment. Edges of buildings in the Minneapolis skyline were used as landmarks and were pointed out to the participants. They were instructed to mark their own position in relation to the buildings with an X on a large piece of newsprint which contained a bird's eye view of the landmark positions. It was found that the subjects were quite poor at this type of estimation. However, training significantly improved their performance. They also performed somewhat better on configurations that were considered to be good in the robot experiments.

Levin and Haber [Levin and Haber 1993] did similar experiments with subjects estimating the distance between two objects. The objects were, in this case, stakes placed at various positions in a grassy field. They showed that the estimate was a linear function of the true distance and the visual angle subtended between the objects. Trials were run with the changes in visual angle caused by alignment, viewing distance and interobject distance. Regardless of cause, the larger visual angles produced larger overestimates of interobject distance.

Waller [Waller 1999; Waller et al. 2000] has run several experiments with human subjects estimating interobject distances in virtual environments. He used both a desktop display and an immersive environment. The number of visual cues were significantly less than those in the outdoor experiments. Texture, for example, was not used since texture and known object size can combine to aid distance estimate. Half the subjects were given error corrective feedback. Waller found that the subjects were able to better estimate exocentric distance than egocentric distance and that estimates were better in the immersive environment. He also found, in agreement with Levin and Haber, that the estimate of interobject distance was increasingly overestimated as the visual angle subtended to the two objects increased.

The experiments described in this paper were modeled after those of Wagner et al. Thus, unlike the interobject distance estimates in the Waller experiments, subjects attempted to locate themselves on a map using visual angle measurements to landmarks in the environment. However, as with the Wagner experiments, although exact distance measurements are unknown, relative distance between landmarks affects the self localization process.

2 Experimental Environment

Subjects sat on a stool at a table on which was located a Wacom graphics tablet. Head movement was allowed, but the subjects were not permitted to leave the stool. They held a pen, with the point placed on the tablet (Figure 1). The room was completely dark to prevent them from seeing objects outside of the virtual environment. They wore a Kaiser ProView XL50 HMD with a 40 degree horizontal field of view, through which they could see the virtual view straight ahead. The environment was rendered with very few visual depth cues. The scene consisted of two wooden fence-like barriers separated by a viewing area in which three wooden poles of equal diameter could be seen (Figure 2). The fence sections extended to the ground and a sufficient distance horizontally and vertically to block all cues which could relate height and ground position to distance. In what was referred to as the “straight-on” view, the outer poles were placed on a line parallel to and 93 meters from the fence. The center pole ranged from approximately 49 to 149 meters from the fence. This changed, as described below, for the skewed views.



Figure 1: Subject during experiment.

Three of the configurations can be seen in Figure 4. A single configuration could be presented with a straight-on viewpoint and with the viewpoint skewed 15, 20 or 30 degrees to the left or right of the center of the configuration. This skewing of the viewpoint resulted in unequal visual angles to the two pairs of landmarks, left/center and right/center. Pole diameter remained the same throughout the experiments, so the further away the pole was, the smaller its diameter. Since the subjects always knew which side of the configuration they were on, this reinforced the view/map relationship. If the center pole was pushed back, for example, it was expected that it would look narrower in the view.

When the subjects looked down in the virtual environment, they saw a virtual tablet (Figure 3). When they moved the real pen on the real tablet, the virtual pen moved accordingly on the virtual tablet. The subjects could press a button on the side of the pen to place a marker at the point on the map where they thought they were located. If they were not satisfied with their answer, they could change it as many times as desired by pressing the button on the pen again. The locations of all of each subject’s answers were recorded. When the subjects were satisfied with their answer, they were shown the next configuration.




Figure 2: The virtual view.



Figure 3: The virtual tablet, visible when the subject looks down.

3 Experiments

The configurations were chosen so that all poles were in the field of view, thus, visible without requiring head turning. Twenty-four configurations were presented to each subject. These consisted of two identical sets of 12 configurations, with one set showing the poles on the map as two dimensional circles and the other set showing them as three dimensional cylinders. All 24 configurations were ordered randomly, but the same order was used for all subjects.

The first experiment was run on 20 subjects, 10 with a scientific background and 10 without. The subjects with a “scientific background” were students with a major in one of the sciences (most were CS, mathematics or physics majors), faculty members in the sciences or members of the Augsburg College IT staff. The “non-science” subjects consisted of students and faculty in the humanities as well as other individuals with little scientific background.

The second experiment was again run on 20 different subjects, 10 with a scientific background and 10 without. The configurations were presented in the same order as in the first experiment, but feedback was given on the first six configurations. The feedback consisted of showing them the true viewpoint on the map after they had made their final decision. The remaining 18 configurations were shown without feedback.




Figure 4: Three of the 24 configurations. Viewpoint is at the star. Actual distances from the horizontal line passing through viewpoint to center poles are 49.55, 99.55 and 149.55 meters, respectively.

4 Results

Results are presented here in three areas: those having to do with the geometry of the configurations and map, those relating to the subjects' prior experience and those relating to the strategies the subjects used.

4.1 Geometric Issues

Performance of both groups directly correlated to the “goodness” measure of the configuration as shown by Sutherland [Sutherland 1993; Sutherland 1994] in the robot navigation work. Performance was better with the center pole pulled forward. Performance became worse as the “single circle” configuration was approached. In short, the better the configurations were for robot localization, the better they were for human localization. This provides some insight into the significance of actual measurement versus perception. The robot sensors made measurement errors, but no perceptual errors were involved. Table 1 shows the results of the first experiment based on “goodness” of configuration. It groups configurations into

3 categories: those with the center landmark closer to the viewpoint, those with the configuration either single circle or straight line and those with the center landmark pulled back.

Configurations	Mean distance to true location
center closer	29
circle or line	62
center back	55

Table 1: “Good” configurations led to better localization. Distance is in meters.

Since the 3D view to 2D map transition has often been described as difficult to make in any localization task, half of the configurations were presented with a flat map and half with 3D poles representing the landmarks. The goal was to test if the 3D map would allow for an easier transition and better localization. There was no improvement with use of the 3D representation. Since this did provide two occurrences of each scenario, the second occurrence, whether two or three dimensional, was checked for improvement in localization in the first experiment. Again, there was no improvement when there was no feedback.

4.2 Prior Experience

Having a background in the sciences significantly improved performance in these localization tasks. Table 2 shows the differences in performance depending on background.

Configurations	Science	Non-Science
center closer	25	34
circle or line	55	63
center back	51	60

Table 2: Mean distance in meters from actual viewpoint for subjects with and without a scientific background.

For both experiments, each subject filled out a pre- and post-test questionnaire. The pre-questionnaire asked if the subjects played video games. There were three possible choices: > 3 hours/week, <= 3 hours/week or not at all. Since 15 of the 20 scientists answered this question in the affirmative (five > 3 hours/week) and the remaining five were all science faculty (who did perform well), and only two of the 20 non-scientists answered in the affirmative, it is impossible to deduce from these experiments how experience in playing video games might affect performance. Those who played video games did perform better, but also had a strong background in the type of scientific thinking which requires making metric observations.

4.3 Strategies

The post-questionnaire asked about strategies employed in the task. Many said that they used the thickness of the poles. The comment “moved head horizontally” was common although only one subject referred to it as motion parallax. A few said that they compared the possible angles on the map to those in the view or that they used triangulation.

In spite of stated strategies, there was a strong tendency to locate on the line of symmetry of the configuration. The subjects with a scientific background were more likely to interpret unequal angle measures in the view (or unequal distances between landmark pairs) as a sign that the viewpoint was skewed.

Multiple answers were allowed and all answers were recorded. The goal was to reveal possible strategies for incrementally improving estimation. These sorts of strategies are often difficult to put into writing. Thirty percent of the results were from multiple estimates. Sixty percent of those had the same value on the first estimate as on the last. This is probably due to the fact that many of the subjects, after deciding on their location, pressed the button on the pen more than once. (This was in spite of the fact that on the first press, a marker appeared at their chosen location.) Twelve percent of the estimates with multiple answers ended at a worse location than they started at while 28% showed improvement. This was certainly not an overall significant improvement. However, allowing the multiple answers made the subjects less anxious about estimating their location, and we intend to allow this in future experiments.

Configurations	Feedback		No feedback	
center closer	11	12	11	25
circle or line	27	21	25	45
center back	20	25	20	53

Table 3: Mean distance in meters from actual viewpoint for subjects with and without feedback. Lateral distance is given by X, straight-on distance by Y.

Feedback improved outcome more than anything else. There is a strong tendency to self locate closer to the configuration than actual viewpoint location. Showing that the actual location was further away on all six trials caused the error in that dimension to decrease significantly. Results are shown in Table 3. However, it was interesting to note that with the improvement in that estimate, the estimate of lateral location became slightly worse. This was especially noticeable with the straight line configurations. It was as if the subjects knew they needed to move back, but it did not seem correct, so they compensated by moving laterally. Figure 5 shows a scatter plot demonstrating this effect.

5 Summary

To summarize, two sets of self localization experiments were performed in an effort to gain insight into spatial judgments in immersive virtual environments. Some parameters were shown to have no effect on self localization. These included representing a map in 3D, seeing a configuration more than once without feedback and being allowed to provide multiple answers. Some parameters clearly influenced self localization. These included the particular landmark configurations used, having a background in the sciences (assumably providing more experience in making metric estimations) and feedback. However, it is not clear if the feedback helped in subsequent estimations or put the subjects into a “move back but I’m not sure why” mode.

Although interobject distances were not being estimated, the overestimation of those distances in the previously cited experiments by Waller and Levin/Haber correspond to the tendency of the subjects to put themselves closer to the configurations than they actually were.

The results of the second experiment (use of feedback) posed an interesting question: What is causing subjects to move their estimated location laterally to compensate for “knowing” that their perception of location is too close? We are currently addressing that issue.




Figure 5: The figure at the top shows a scatter plot of estimated locations in the first experiment (no feedback) for the straight line configuration shown in red. The true location is at the blue star. The figure below shows performance in the second experiment (feedback on first six configurations). Note how the improvement in estimate of distance back from the configuration is accompanied by lateral movement.

References

- LEVIN, C. A., AND HABER, R. N. 1993. Visual angle as a determinant of perceived interobject distance. *Perception and Psychophysics* 54, 2, 250–259.
- SUTHERLAND, K. T. 1993. Landmark selection for accurate navigation. In *Proc. DARPA Image Understanding Workshop*, 485–490.
- SUTHERLAND, K. T. 1994. The stability of geometric inference in location determination. Tech. Rep. UUCS-94-021, University of Utah, Salt Lake City, July.
- WAGNER, D., PICK JR., H. L., THOMPSON, W. B., YONAS, A., AND WEGESIN, D. 1994. Localization by visual angle. Tech. rep. Unpublished.
- WALLER, D., LOOMIS, J. M., GOLLEDGE, R. G., AND BEALL, A. C. 2000. Place learning in humans: The role of distance and direction information. *Spatial Cognition and Computation* 2, 333–354.
- WALLER, D. 1999. Factors affecting the perception of interobject distances in virtual environments. *Presense: Teleoperators and Virtual Environments* 8, 657–670.