The Hand is Slower than the Eye:
A quantitative exploration of visual dominance over proprioception

Eric Burns
Mary C. Whitton

Sharif Razzaque
Matthew R. McCallus

Abigail T. Panter
Frederick P. Brooks, Jr

University of North Carolina, Chapel Hill

Figure 1. A participant who believes he is aiming at the virtual game board directly in front of him.

ABSTRACT

Without force feedback, a head-mounted display user's avatar may penetrate virtual objects. Some virtual environment designers prevent visual interpenetration, making the assumption that prevention improves user experience. However, preventing visual avatar interpenetration causes discrepancy between visual and proprioceptive cues. We investigated users' detection thresholds for visual interpenetration (the depth at which they see that two objects have interpenetrated) and sensory discrepancy (the displacement at which they notice mismatched visual and proprioceptive cues). We found that users are much less sensitive to visual-proprioceptive conflict than they are to visual interpenetration. We present our plan for using this result to create a better technique for dealing with virtual object penetration.

burns@cs.unc.edu sharif@cs.unc.edu
panter@unc.edu whitton@unc.edu
mccallus@cs.unc.edu brooks@cs.unc.edu

CR Categories: H.5.2 [Information Interfaces and Presentation]: User Interfaces; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality
Keywords: sensory discrepancy, detection thresholds, perception

1 INTRODUCTION

To be realistic, a virtual environment (VE) must induce perceptions consistent with those induced by the real world [1]. Perfectly mimicking real-world stimuli would be one way to achieve this goal. However, previous research suggests that easier ways exist.

Razzaque, Kohn, and Whitton [2] exploited the fact that our perceptual systems weight visual cues more heavily than vestibular cues. By presenting a distorted visual scene when users turned their bodies, they made users experience a rotation different from that which was actually occurring in the real world. They used this technique to address the problem of limited tracker space.

Razzaque, et al. introduced a distortion that was unlike the real world to make users perceive the VE as more like the real world (because they were not restricted in where they could walk). This work has inspired us to explore other aspects of VE that can be improved using this same pattern. Specifically, we have looked at avatars in head-mounted display (HMD) VEs. Heeter [3] found that having an avatar increases a user's sense of presence. However, it is impractical for many VEs to offer force feedback.

Without force feedback, a user may move his body such that it occupies the real space that corresponds to the location of a virtual object in the VE. The user would then see his avatar penetrate the virtual object (Figure 2), which is inconsistent with real-world experience and may cause a break in presence [4].

Figure 2. Without force feedback, a user may move his body such that his avatar penetrates virtual objects.

Users easily detect this visual interpenetration. The seemingly logical conclusion is that a VE would be improved by preventing the user's avatar from penetrating virtual objects. But this conclusion may not be valid. Take the example of a user who has an avatar hand in the VE.

Visual penetration can be prevented because the visual scene is completely generated by the simulation. However, without force feedback, nothing can prevent the user from moving his hand such that it occupies the space corresponding to the location of a virtual object. Since the user can move his real hand to places that his avatar hand cannot go, the positions of the user's real and avatar hands can be decoupled. When this happens, the user's visual cues (the seen position of his avatar hand) and proprioceptive cues (the felt position of his real hand) are discrepant (Figure 3).

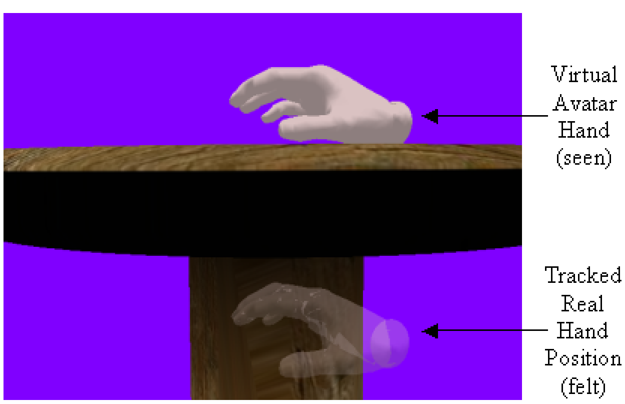

Figure 3. Preventing avatar interpenetration causes the seen position of the hand to differ from the felt position of the hand.

We are only justified in stopping visual interpenetrations if this sensory discrepancy is less noticeable or less detrimental than the original visual interpenetration. We conducted a user study to address the former – whether users are more sensitive to visual interpenetration or visual-proprioceptive conflict.

We found that users are much more sensitive to visual interpenetration than to visual-proprioceptive conflict. More interesting than the comparison is *how* insensitive users are to visual-proprioceptive conflict. We report our results as well as a new technique for dealing with virtual object penetration.

2 Background

Psychologists have known for decades that vision usually dominates proprioception when the two disagree [5]. Many researchers have undertaken to explain intersensory dominance [6] – [9]. Warren and Cleaves [9] found that under large amounts of discrepancy, the dominance is not complete. Van Beers, Wolpert, and Haggard [8] even showed that in certain situations, proprioception can dominate vision. However, the participants in Warren and Cleaves and van Beers et al. were unaware of the sensory discrepancy. Wann and Ibrahim [10] and Brown [11] found that without visual feedback, perceived limb position drifts over time. We are unaware of any research to determine the extent of discrepancy possible before participants notice. We are also unaware of any work to determine the extent of visual penetration possible before a participant notices.

3 Study Design

Forty right-handed introductory psychology students (19 males and 21 females) participated in this study. They were given class credit for their participation.

The study consisted of three parts. Part I measured reaction time. Parts II and III used a partial *method of limits* to detect perceptual thresholds for visual-proprioceptive discrepancy and visual interpenetration respectively. We used an ascending series (starting with no stimulus and increasing the stimulus until the user reports that it is perceived – that is, the perceptual threshold has been reached) for both. All participants completed Part I first, but the order of Parts II and III were assigned randomly.

3.1 Part I – Reaction Time

In this study, we were interested in users' detection times, but we could only measure their report times. If the report time is the detection time plus the time necessary for the participant to react, we can estimate the participant's detection time by subtracting his average reaction time from his report time. To make this possible, we measured the participants' reaction times in the first part of the study.

Each participant sat in front of a black computer screen and held a joystick in his right hand. At randomly chosen intervals the screen turned white, at which point participants clicked the button on the joystick as quickly as possible. The time between the screen changing white and the button press was recorded as their reaction time. Participants performed this task 45 times.

We made the assumption that adding this task to the beginning of the study would not significantly affect users' performance on the subsequent tasks. We believe this is a valid assumption because the task was dissimilar enough from those that followed that it was unlikely to result in a significant training effect, and it was short enough that it was unlikely to result in fatigue effects.

3.2 Part II – Virtual Simon®

This part of the study measured participants' detection thresholds for visual and proprioceptive discrepancy. Each participant wore a Virtual Research Systems V8 HMD and held a joystick in his right hand. Both the head and hand were tracked by 3rdTech Hiball sensors. The participant sat in a chair (Figure 1) and was visually immersed in a virtual room that had four large colored panels on a wall in front of him. He had an avatar hand with a remote control in it (Figure 4).

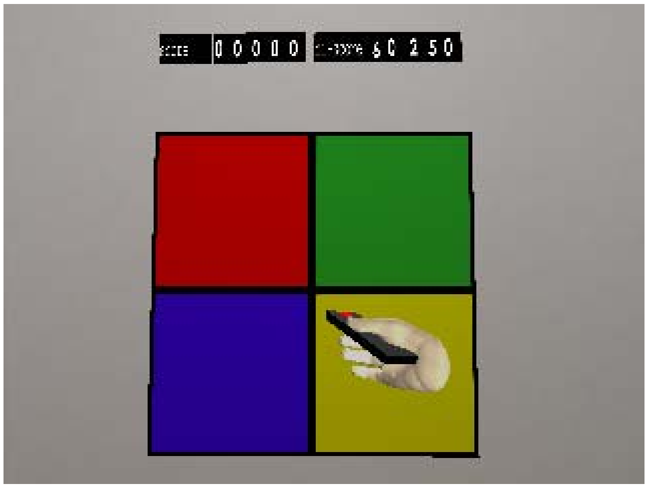

Figure 4. A view from within the virtual room, with the Simon game board on the wall. The user's avatar hand, holding a TV-like remote control, is in the foreground.

We instructed participants to play a game similar to Hasbro's Simon® in the VE. The participants watched as the panels lit up in a random sequence of five colors. They then repeated the pattern by aiming at each panel with the right hand and clicking the button on the joystick. The simulation calculated a ray from the head through the center of the virtual avatar hand and found the point that it intersected the wall. If that point fell within the bounds of one of the colored panels, that panel was selected. After the participant completed each sequence correctly or made an error, a new sequence began. To keep participants engaged in the game, we scored them on their performance. The score was displayed on the wall overtop the colored panels, together with a top score over all participants to date (like a typical arcade game).

Before the game began, participants were told that our study was about perception and performance in a VE and therefore, it was very important for us to know if they noticed anything odd about the VE experience. They were told to report anything odd immediately by holding down the button on the joystick for five seconds. We then gave them three examples of events they would want to report: the game stopping, the computer display having problems, or the virtual hand seeming to have drifted away from where their real hand actually was.

The game was divided into nine trials. The first trial measured the sensory discrepancy detection threshold when a user *was not* making its detection an attentional priority. After a random length of time that averaged 25 seconds and was geometrically distributed, the participant's avatar hand was made to drift from the real hand position. The hand drifted counter-clockwise as viewed from above along a circle that had its center at the participant's estimated shoulder position (a fixed offset from the head-tracker) (Figure 5). This drift is named the *left* drift condition because the avatar hand moves left across the user's field of view.

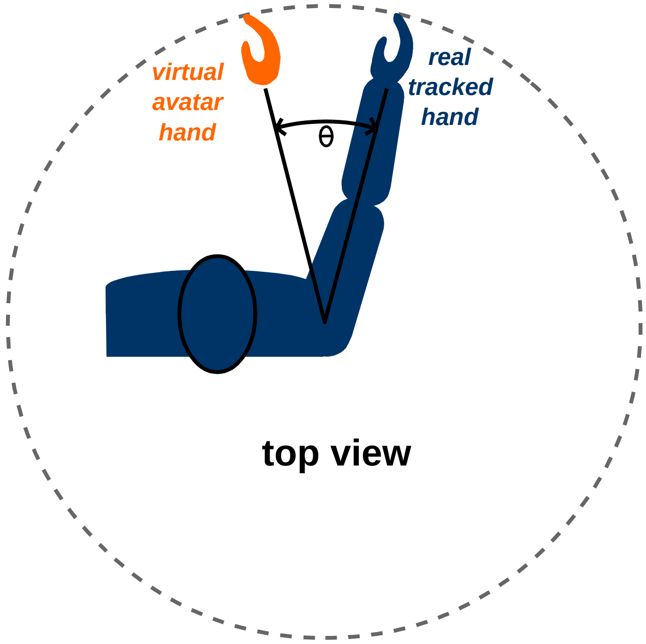

Figure 5. After a random length of time, the participant's virtual avatar hand began to drift counterclockwise about the estimated shoulder position (as viewed from above).

Since we were interested in the maximum distance between the avatar and real hands before a user noticed, it was important that it was really the extent of the drift that they detected and not the motion of the avatar hand itself. Therefore, it was very important that the drift be executed in such a way that we were sure it was imperceptible. Pre-study piloting showed that participants could detect the motion of even a very slow drift if they held their hands completely still and watched for it. Therefore, during the study, the avatar hand only drifted if the user's hand was moving faster than 5 cm/s. When the user's hand was moving above this threshold, the avatar hand drifted 0.458 degrees/s (5 mm/s for someone with a 63.5 cm long arm). Both values were chosen because in pre-study piloting no one was able to detect the motion of the drift.

If the drift reached 60 degrees without the participant reporting anything odd, the trial was stopped, and the user was asked if he had noticed anything odd. If he said, "No," he was told that there was something odd and was asked to guess what it was. If he did not guess correctly, he was told that the hand had drifted and was then asked again if he had noticed. Results are presented in section four.

In four of the remaining eight trials, the hand drifted left, right, up, and down. In the other four trials, the hand did not drift. The drift condition order was selected from an 8x8 balanced Latin square matrix. The experimenter told participants that in each remaining trial the hand would have a 50 percent chance of drifting. He instructed them to report the drift as soon as they noticed it and also report which direction the hand drifted. He told them that it was much more important to report the drift immediately upon noticing it than it was to get the direction correct. He then told them they would be rewarded with points for correctly identifying drift regardless of whether they chose the correct direction, but would be penalized the same number of points for a false alarm. These eight trials measured the sensory discrepancy detection threshold when a user *was* making its detection an attentional priority.

3.3 Part III – Penetration Depth

This part of the study measured each participant's detection threshold for visual interpenetration. Participants wore the same HMD and held the same joystick as in Part II. In this part of the study, the movement of the user's real hand did not control the movement of the virtual hand.

Participants were assigned randomly to one of two conditions. In one condition (the vertical motion condition), participants viewed a virtual hand holding a cylinder above a wood-textured tabletop (Figure 6). At the beginning of each trial the participant clicked the button and the hand began to move down toward the tabletop in front of him. The hand was placed so that the point it would impact the tabletop was at a 45-degree angle below horizontal. Participants were told that in each trial the hand had a 50 percent chance of penetrating the table. If the hand penetrated the table, they were instructed to click the button as soon as they noticed. The speed of the hand was varied so that participants could not solely use time to judge when the hand would penetrate the table. Participants repeated this task 40 times. This condition mimics a typical scenario that would occur if a person were seated at a table and placed his hand on top of it with his arm outstretched. This condition corresponds well to the sensory discrepancy task in Part II because the motion of the hand corresponds exactly to the up drift condition.

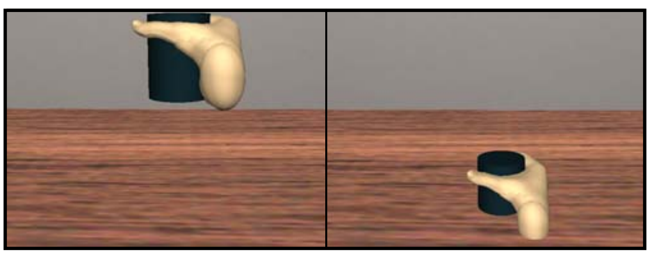

Figure 6. In the vertical motion condition, participants viewed a hand holding a cylinder above a wood-textured tabletop. Left – the starting position of the hand; Right – the hand after penetrating the table 2 cm

We realized that viewing the hand penetration from different angles and with different backgrounds could affect the difficulty of the task. For example, a viewing angle perpendicular to the motion of the hand would make detection easiest because the gap between the objects would be directly visible. Conversely, a viewing angle parallel to the motion of the hand would make detection the most difficult because the user would have to rely on depth cues to detect the penetration, and the point of contact would be obscured by the hand itself until it became extreme (Figure 7). Since we have hypothesized that sensory discrepancy is harder to detect than visual interpenetration, we were afraid that our choice of visual penetration condition would be biased toward making it easy. Therefore, we decided to add another condition that we named the *horizontal motion* condition. In this condition participants viewed the same virtual hand 20 cm in front of a featureless wall that offered minimal depth cues (Figure 8). At the beginning of each trial, participants clicked the button and the hand began moving away from them toward the wall. Participants were given the same instructions as in the vertical motion condition.

Figure 7. It is easier to detect the collision of a ball with the ground when viewed from the side (perpendicular to the direction of motion), as seen in the left picture, than when viewed from above (parallel to the direction of motion) as seen in the right picture.

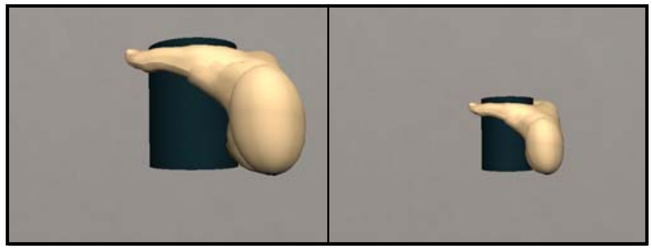

Figure 8. In the horizontal motion condition participants viewed a hand holding a cylinder in front of a featureless wall. Left – the starting position of the hand; Right – the hand after penetrating the wall 2 cm

The penetration and hand speed orders were selected from independent 40x40 balanced Latin square matrices. In both conditions participants were restricted from moving their heads to view the hand from a different angle but were free to look around the room and gather depth cues from the other walls.

4 RESULTS

We lost six sensory discrepancy values due to software malfunctions, sixteen due to false alarms on trials where the hand would have drifted (when the participant reported drift before it began), and sixteen due to time constraints. These missing values left 19 out of 40 participants who had complete data.

4.1 Sensory Discrepancy Drift Directions

A repeated measures ANOVA failed to find a significant difference in sensory discrepancy thresholds between the four drift directions for the 19 participants with complete data ($F(3, 54) = 0.8$, $p < 0.498$). However, this analysis ignores the possibility that participants with missing data vary systematically with respect to the participants with complete data. Since participants at the two extremes of performance (under-responders, who took a long time to report and may have run out of time before completing the experiment; and over-responders, who may have reported drift before it actually occurred, registering a false alarm rather than a valid data point) are the most likely to have missing data, we are unwilling to claim that they do not vary systematically from participants with complete data.

In order to include the effect of participants with missing data, we used a multiple imputation method to generate thirty complete datasets that preserved the means and variances of the observed data. We then performed a two-tailed t-test on the six direction pairs for each dataset and combined the results to produce the statistics shown in Table 1. None of the pairs returned a significant difference.

Table 1. Results of the two-tailed t-test for each direction pair on the multiply imputed dataset of sensory discrepancy thresholds

Direction pair	t	p
left / right	1.92	0.0627
left / up	1.05	0.298
left / down	0.34	0.736
right / up	-1.03	0.311
right / down	-0.803	0.427
up / down	-0.29	0.774

Since neither our test on the participants with complete data nor our test on the imputed datasets returned any significant differences, we would like to conclude that there is likely no difference in the detection threshold for the different directions. However, this conclusion is flawed. First, it is improper to conclude that since we did not find a significant difference, one does not exist (especially with the left/right pair approaching significance). Second, both our tests have orthogonal issues that call their credibility into question. As already mentioned, the test on participants with complete data excludes participants that may vary systematically. Our multiple imputation method probably gives us a better picture of reality, but it is questionable because 20 percent of our data are missing. It is commonly (though arbitrarily) accepted that imputation is an effective method of accounting for missing data up to five percent of the total dataset. We have stretched the method far beyond that.

For all of these reasons, we can make no solid conclusions about how detection thresholds for visual-proprioceptive discrepancy differ with direction of arm motion. However, despite these difficulties, we will assume the null hypothesis and treat the different directions of drift as four measurements of the same sensory discrepancy threshold. As the next section shows, this unfounded assumption is insignificant compared to the differences measured between tasks and it simplifies our data analysis considerably.

4.2 Detection Thresholds

The mean detection thresholds for visual-proprioceptive discrepancy and visual interpenetration are shown in Figure 9. These values represent the estimated stimulus levels at the time of detection, calculated from report times and reaction times (mean reaction time = 260 ms, standard deviation = 20 ms) as follows:

detection threshold = stimulus at time of report
– (reaction time)(speed of hand)

We discarded false alarms prior to calculating the mean detection thresholds.

Our estimate of the sensory discrepancy threshold is conservative for two reasons. First, the virtual avatar hand was assumed to be moving during the duration of the participant's reaction time. If the user held his hand still or removed it from his field of view, the hand would not have moved during this time and the value subtracted from the mean discrepancy at the time of report would be too large. Second, the mean detection threshold ignores the false alarm rate of the participants. Figure 10 shows mean detection thresholds as a function of the number of false alarms reported by the participant. A linear regression of mean detection threshold on number of false alarms yielded a statistically significant downward trend (intercept = 0.2268m, slope = -0.0217m, $F(1, 31) = 8.68$, $p < 0.0061$). This result leads to the conclusion that the participants with the lowest thresholds also had the highest number of false alarms. This conclusion suggests that these participants were not consistently able to discriminate sensory discrepancy from a lack of sensory discrepancy and were most likely guessing.

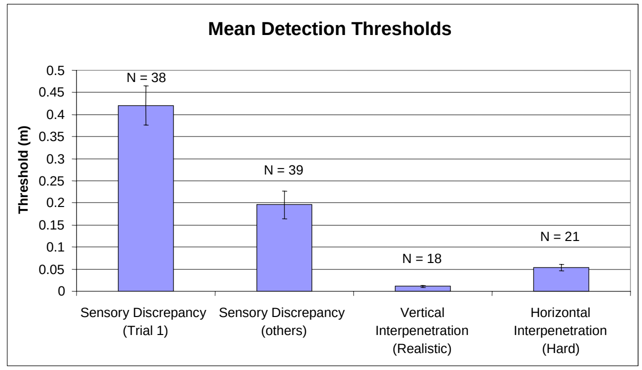

Figure 9. Mean detection thresholds for visual-proprioceptive discrepancy and visual interpenetration – Bars represent a 95 percent confidence interval for the mean (1.96 * standard error).

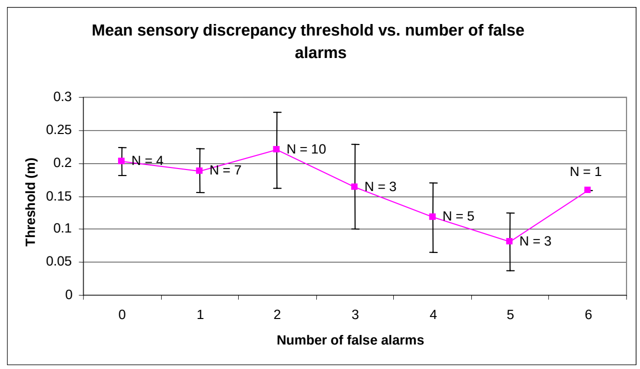

Figure 10. Mean detection thresholds as a function of the number of false alarms reported by the participant – Bars represent a 95 percent confidence interval for the mean and N values represent the number of participants with each false alarm rate. Seven participants were excluded because time constraints did not allow them to finish all trials.

Analyzing receiver-operator curves to determine the discriminability of the visual-proprioceptive discrepancy would address participants' guessing and bias. However, the number of trials required to generate such curves was impractical for this study. We instead used the data from all participants without regard to false alarm rates to estimate a mean detection threshold. Our resulting estimate is conservative because the data from participants with high false alarm rates artificially lowers the mean.

We used MANOVA to analyze these threshold data because we did not wish to assume that detection thresholds for the two tasks would have equal variances. Despite our conservative estimate of sensory discrepancy thresholds, the analysis showed a significant difference between sensory discrepancy thresholds on the last eight trials and visual interpenetration thresholds for both the vertical motion condition ($F(1, 17) = 61.74$, $p < 0.0001$) and the horizontal motion condition ($F(1, 20) = 322.23$, $p < 0.0001$).

4.3 Self Report of Task Difficulty

On an exit questionnaire, participants were asked to rate the difficulty of detecting the sensory discrepancy and visual interpenetration on a scale from one to seven. An ordered multinomial regression showed that on the exit questionnaire participants rated the drift task significantly harder than the visual interpenetration task with $\chi^2(1) = 62.7$, $p < 0.0001$ (Figure 11).

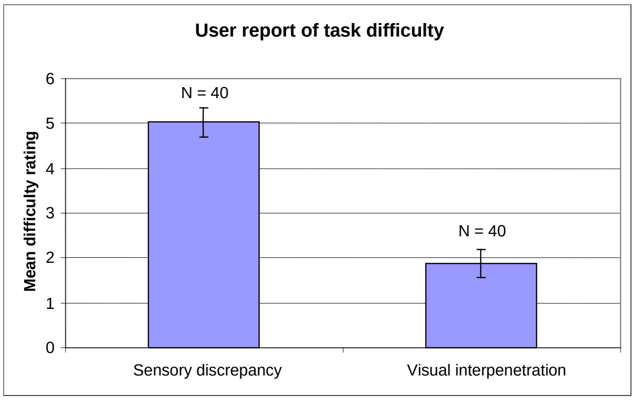

Figure 11. User report of task difficulty on a scale of 1 to 7 (with 1 as easiest and 7 as hardest) – Bars represent a 95 percent confidence interval for the mean.

4.4 Sensory Discrepancy Thresholds as a Function of Expectation

A repeated measures t-test showed a significant difference between the sensory discrepancy threshold on the first trial (when participants were not expecting hand drift) and the average thresholds from the last eight trials (when they were expecting hand drift) with $t(33) = 9.008$, $p < 0.0001$. However, this result is inconclusive due to the systematic underestimation of the detection threshold on the last eight trials due to over-reporting (having a high false alarm rate). A linear regression of first trial detection threshold on false alarm rate failed to show the same trend as the thresholds for the last eight trials. Therefore, we are not able to assume that both values are subject to the same underestimation.

However, the first trial's detection threshold is underestimated for a different reason. Seventeen participants did not report an odd event on the first trial. The trial ended after their virtual avatar hand reached an offset from their real hand of 60 degrees. Of these participants, four said they did not notice their hand had drifted at all. Eight could not guess what was odd about their experience when they were told that we had introduced a manipulation to see if they would notice. However, when asked if they noticed that the virtual hand had drifted they said they did notice. One of these participants volunteered his understanding of where his real hand was in relation to his avatar hand, but did so incorrectly. The other five non-reporting participants immediately mentioned the hand drift when asked if anything was odd about their experience, even though they had not reported it. These five participants were not included in the statistics for trial one. All others were included. In addition to the 17 participants who never reported on the first trial, eight participants reported some other odd occurrence before they noticed the hand had drifted. As a result, we can only be sure that the reported value for 14 out of 39 participants (because one participant's first trial was lost due to an equipment malfunction) represents an actual detection threshold. The other reported values represent a lower bound on the participants' thresholds. We cannot infer how low this estimate is.

4.5 Caveat About the Ascending Series

Though several factors are at work that underestimate the sensory discrepancy thresholds, it is commonly accepted that using an ascending series (as we have in this study) will *overestimate* detection thresholds. A true method of limits design balances the ascending series with a descending series (which underestimates detection thresholds) and combines the results to find a more accurate estimate of the detection threshold. However, our aim is to determine how large these stimuli can *grow* before they are noticed. Therefore, we feel that the ascending series design is most appropriate and does not overestimate the threshold we wished to determine.

4.6 Performance Effects

For each trial of the Simon game, we calculated the participant's score per second (total score for that trial divided by the total number of seconds in that trial). A repeated measures t-test showed a significant difference between trials during which the hand drifted and those in which it did not with $t(39) = 3.18$, $p < 0.0029$ (Figure 12).

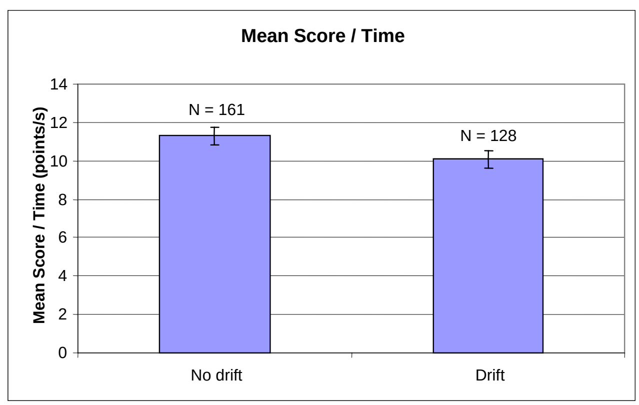

Figure 12. Mean score / time in trials where the hand drifted as compared to trials in which it did not – Bars represent a 95 percent confidence interval for the mean.

5 IMPLICATIONS FOR VEs

The results supported our hypothesis — humans are more sensitive to visual interpenetration than visual-proprioceptive conflict. This result suggests that separating the real hand and visual avatar hand to prevent visual interpenetration is beneficial. However, two other details are needed to validly draw this conclusion. First, we must be certain that when visual-proprioceptive conflict in hand position grows large enough to be noticed by the user, it is less detrimental to his experience than the corresponding visual interpenetration, or that the discrepancy will be noticed infrequently enough to warrant its introduction. Second, we must be certain that, on average, the visual-proprioceptive discrepancy in hand *velocity* is also less detrimental than visual interpenetration. We intend to perform future studies to investigate these issues.

We believe that preventing visual interpenetration will indeed prove helpful. However, decoupling the position of the avatar hand from the position of the user's real hand introduces new concerns. Prior to such separation the position and orientation of the visual avatar hand are set directly from the tracked position and orientation of the real hand. However, once we have separated the two, it is unclear how to best manage the position and orientation of the visual avatar hand so as to continually present the most perceptually plausible experience to the user.

Zachmann and Rettig [12] addressed this issue in their paper on interaction in virtual assembly simulation. Though the paper's main focus lay elsewhere, they proposed two possible methods for determining a visual avatar hand's position after being separated from the user's real hand.

The first is the *rubber band* method. In this method, the visual avatar hand maintains a position and orientation as close as possible to the participant's real hand. This method will cause the visual avatar hand to appear to stick to the surface of a virtual object as the user draws his hand out of the object (Figure 13). This method minimizes the discrepancy between the *position* of the visual and proprioceptive hands at the expense of creating a discrepancy between the *motion* of the visual and proprioceptive hands.

Their second method is the *incremental motion* method. In this method the real hand's change in position is calculated at every time step. This displacement is applied to the visual avatar hand as well (Figure 14). This method minimizes the discrepancy between the *motion* of the visual and proprioceptive hands at the expense of maintaining the discrepancy between the *positions* of the visual and proprioceptive hands.

Neither of these techniques is ideal. Our pre-study pilot showed that users are very sensitive to discrepant motion. Therefore, the rubber band method is inadequate. Our data suggest that the incremental motion model is inadequate, as well. Participants' Simon game scores were significantly lower on trials with sensory discrepancy than on those without. Seven participants commented that they felt that sometimes they would aim at one panel and another would light up. We believe this phenomenon is due to incomplete dominance as described by Warren and Cleaves [9]. Though participants were not consciously aware that their visual and proprioceptive cues were discrepant, the discrepancy affected their perceived hand positions. They used their perceived hand positions to aim at the desired panel, but since a ray through the virtual avatar hand made the selection, sometimes a different panel was selected than the one they desired (Figure 15). Six participants commented in the exit interview that they even used aiming errors as a cue that the virtual avatar hand had drifted. By maintaining sensory discrepancy in position, the incremental motion model continually induces performance errors.

6 Credible Avatar Limb Motion

Both methods of managing the virtual avatar hand achieve the minimum discrepancy in one variable only. They can be described as opposite ends of a continuum (Figure 16). We propose a new technique called *Credible Avatar Limb Motion* (CALM) that blends these two extremes to create the minimum sensory discrepancy as a function of both displacement and velocity. In this method, the offset between the virtual avatar hand and the user's real hand is maintained as in the incremental motion method. However, when the user moves in such a way that it becomes possible to bring the two hands closer together, the virtual avatar hand is slowly moved toward the user's real hand until they meet. This technique would trade sensory discrepancy in position for velocity and vice versa in order to ensure that neither discrepancy becomes intolerably large.

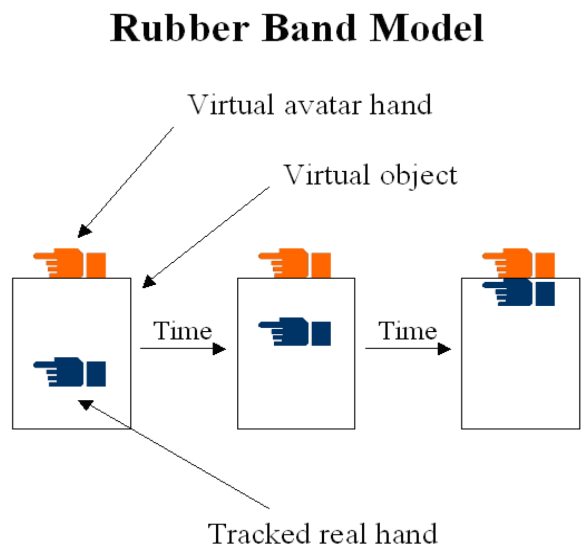

Figure 13. Because the rubber band method minimizes the distance between virtual avatar hand and the user's real hand, the virtual avatar hand may appear to stick to surfaces.

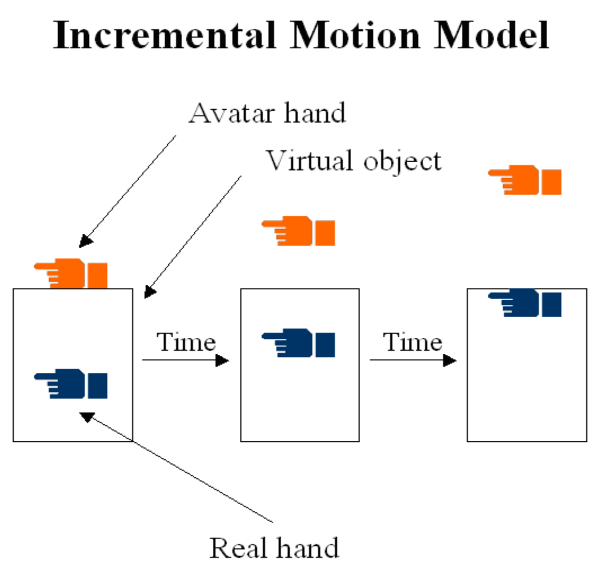

Figure 14. The incremental motion model minimizes the discrepancy between the motion of the visual and proprioceptive hands but maintains the discrepancy in position.

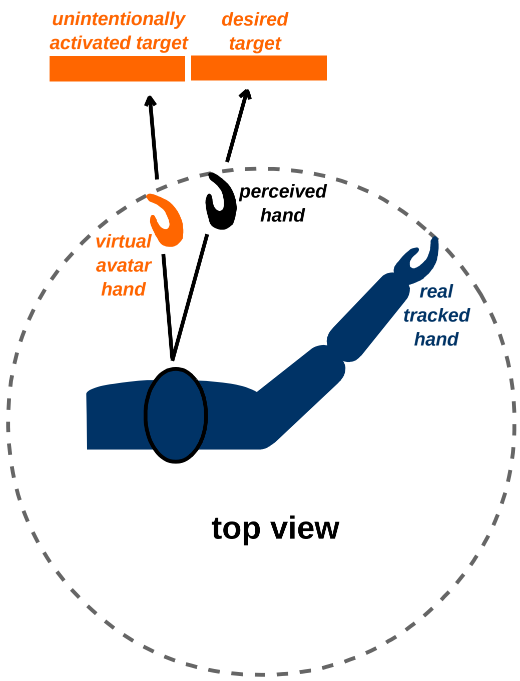

Figure 15. Large sensory discrepancies can lead to incomplete dominance of vision over proprioception. This incomplete dominance caused performance errors as participants aimed with their perceived hand position, which occasionally differed from their virtual avatar hand position by enough to activate a panel other than the one desired.

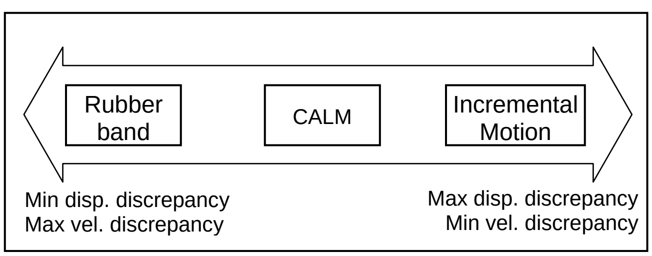

Figure 16. The rubber band and incremental motion models represent opposite ends of a continuum of displacement and velocity discrepancy. The ideal solution to the problem is somewhere between them on that continuum.

It remains unclear how quickly or in what manner to move the virtual avatar hand toward the user's real hand. We plan to investigate these parameters and whether this method of dealing with virtual avatar hand position has any advantage over the two extremes.

Independent of the to-be-determined effectiveness of CALM, we have found that users notice visual interpenetration more easily than visual-proprioceptive discrepancy in displacement. This discovery is the first step in determining whether VE designers should prevent users' avatars from visually penetrating virtual objects.

7 Acknowledgments

Support for this research was provided by the Office of Naval Research, the NIH National Institute for Biomedical Imaging and Bioengineering, and the NIH National Institute for Research Resources. We thank Chris VanderKnyff, Luv Kohli, Paul Zimmons, and Chris Oates for their assistance with the user study. We are grateful for advice from Bernard Adelstein at the NASA Ames Research Center and Leonard McMillan and Chris Wiesen at UNC.

References

[1] B. Witmer and M. Singer, "Measuring presence in virtual environments: A presence questionnaire." *Presence: Teleoperators and Virtual Environments*, vol. 7, no. 3, 1998, pp. 225-240.

[2] S. Razzaque, Z. Kohn, and M. Whitton, "Redirected walking." *Proceedings of Eurographics 2001*, Sep. 2001.

[3] C. Heeter, "Being there: The subjective experience of presence." *Presence: Teleoperators and Virtual Environments*, vol. 1, no. 2, 1992, pp. 262-271.

[4] M. Slater and A. Steed, "A Virtual Presence Counter," *Presence: Teleoperators and Virtual Environments*, vol. 9, pp. 413-434, 2000.

[5] J. Gibson, "Adaptation, after-effect and contrast in the perception of curved lines." *Journal of Experimental Psychology*, vol. 16, no. 1, 1993, pp. 1-31.

[6] M.I. Posner, M.J. Nissen, and R.M. Klein, "Visual dominance: An information-processing account of its origins and significance." *Psychological Review*, vol. 83, no. 2, 1976, pp. 157-171.

[7] R. van Beers, A. Sittig, and J. Denier van der Gon, "Integration of proprioceptive and visual position-information: An experimentally supported model." *Journal of Neurophysiology.*, vol. 81, 1999, pp. 1355-1364.

[8] R. van Beers, D. Wolpert, and P. Haggard, "When feeling is more important than seeing in sensorimotor adaptation." *Current Biology*, vol. 12, 2002, pp. 834-837.

[9] D. H. Warren, and W.T. Cleaves, "Visual-proprioceptive interaction under large amounts of conflict." *Journal of Experimental Psychology*, vol. 90, 1971, pp.206-214.

[10] J.P. Wann, and S.F. Ibrahim, "Does limb proprioception drift?" *Experimental Brain Research*, vol. 91, 1992, pp. 162-166.

[11] L.E. Brown, *Limb position drift and its implications for movement and position control*, doctoral dissertation, Dept. Psychology, The Pennsylvania State University, 2003.

[12] G. Zachmann and A. Rettig, "Natural and robust interaction in virtual assembly simulation." *Eighth ISPE International Conference on Concurrent Engineering: Research and Applications (ISPE/CE2001)*, July 2001.