

Reaching Movements to Augmented and Graphic Objects in Virtual Environments

Andrea H. Mason, Masuma A. Walji, Elaine J. Lee and Christine L. MacKenzie

School of Kinesiology

Simon Fraser University

Burnaby, B.C. V5A 1S6, Canada

+1 604 291-5794

{ahm, mwalji, ejlee, christine_mackenzie}@sfu.ca

ABSTRACT

This work explores how the availability of visual and haptic feedback affects the kinematics of reaching performance in a tabletop virtual environment. Eight subjects performed reach-to-grasp movements toward target objects of various sizes in conditions where visual and haptic feedback were either present or absent. It was found that movement time was slower when visual feedback of the moving limb was not available. Further MT varied systematically with target size when haptic feedback was available (i.e. augmented targets), and thus followed Fitts' law. However, movement times were constant regardless of target size when haptic feedback was removed. In depth analysis of the reaching kinematics revealed that subjects spent longer decelerating toward smaller targets in conditions where haptic feedback was available. In contrast, deceleration time was constant when haptic feedback was absent. These results suggest that visual feedback about the moving limb and veridical haptic feedback about object contact are extremely important for humans to effectively work in virtual environments.

KEYWORDS

Augmented reality, kinematic data, object manipulation, haptic feedback, visual feedback, sensory information, human performance, Fitts' law, empirical data, interaction

INTRODUCTION

Object manipulation is a fundamental operation in both natural human movement and human computer interaction (HCI). By taking advantage of the human ability to use our hands to acquire and manipulate objects with ease, designers can construct interactive virtual and augmented environments that will be seamlessly and effectively used [18]. However, designers must also consider that our ability

to manipulate objects with ease is strongly related to the sources of sensory information that we gather prior to and after contact with objects [8]. Specifically visual and haptic feedback are key sources of sensory information used when acquiring and manipulating objects. Unfortunately, incorporating rich interactive graphics and haptic feedback in virtual environments is costly both in terms of computing cycles, and equipment purchases. Thus, it is important to determine whether the cost of implementing these sources of feedback can be justified by performance improvements. We describe in this paper an experiment performed to investigate the effects of removing haptic and visual feedback when subjects use their hands to acquire objects in a virtual environment.

Target acquisition and haptic feedback

Much of the research to date on target acquisition in computer generated environments has focused on pointing or aiming movements to targets of various sizes and amplitudes using input devices such as a mouse, trackball or tablet in a standard desktop configuration [11]. Consistent with Fitts' law, it has generally been concluded that movement time increases with increases in index of difficulty [2].

With modern computer systems such as virtual or augmented environments it is possible to achieve multidimensional input using the whole hand as the object manipulation device. In studies where the hand has been used as the manipulation device for aiming to targets in both desktop and virtual environments, movement times have also been found to conform with Fitts' law [3,12]. However in these studies, subjects used their fingers as pointers to planar targets on the table surface and thus haptic feedback was always available at target contact. In the current paper, we are interested in understanding how the absence of haptic feedback at target contact affects movement times and the ability to generalize Fitts' law.

Within the study of human performance in virtual environments, recent research has shown that haptic feedback not only provides realism in virtual environments [4], but also enhances human performance [1,6,20]. Wang and MacKenzie [20] performed a study in which subjects

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCHI'01, March 31-April 4, 2001, Seattle, WA, USA.

Copyright 2001 ACM 1-58113-327-8/01/0003...\$5.00.

moved an object in hand to dock it with a 3-dimensional wireframe graphic cube. In some conditions, the physical table on which the target position was located was present, while in other conditions it was removed. Thus, haptic feedback at object contact with the table surface was manipulated. The task completion time was dramatically increased when the tabletop was absent. However, regardless of whether haptic feedback was available or not, movement time results always followed Fitts' law. Also, Linderman, Sibert and Hahn compared human performance when docking a graphic object to either a 'floating' graphic panel or to a panel that was augmented by a physical paddle[6]. Again, these authors reported that in conditions where haptic feedback was received, subjects were faster at docking the object than in conditions with no haptic feedback. Finally, Arsenaault & Ware reported movement time advantages of 12% when haptic feedback was available at target contact in a Fitts' aiming task within a virtual environment than when target contact was signaled visually [1]. Thus, for object aiming and docking tasks, we have evidence that haptic feedback does improve performance in terms of decreased movement time. We also have evidence that regardless of whether or not subjects receive haptic feedback, Fitts' law holds true.

A notable difference between the experiments conducted in [1,6,20] and the current experiment, is that subjects transported an object already in hand to aim to or dock with a target. However, in the current experiment we are specifically interested in understanding what role haptic feedback plays when subjects *acquire* objects into grasp. When interacting with objects in real world situations, we expect that when we make contact with an object we will receive haptic feedback about the object's shape, texture, and mass [5]. However, with evolving computer technologies, we are beginning to interact with objects that exist only as graphic representations. Thus, do the same laws hold for these 'virtual' interactions when expected feedback is not always obtained? Will the same movement time benefits be seen that were shown in [1,6,20], and will Fitts' law still hold when subjects reach to grasp a completely virtual object?

Target acquisition and visual feedback

In the current experiment, we are also interested in investigating how visual feedback facilitates target acquisition movements. Visual information is extremely important for the performance of many motor activities. It can provide information not only about object properties such as size and orientation, but also about the movement of one's own limbs within the environment. Under normal visual control, target acquisition movements are made quickly, and accurately [16]. However, due to limited computing power, it is not always possible in virtual environments to provide subjects with rich graphic feedback about the objects and relationship between the environment and their moving limbs. It was shown that when vision of the moving limb was removed, errors occurred in the terminal location of aiming movements in natural environments [16].

Furthermore, in a desktop virtual environment, it was shown that subjects took longer to make aiming movements toward computer generated targets when a graphic representation of the finger was not available compared to when the finger was represented by a graphical pointer [12]. Thus visual feedback or a graphic representation of the movement of one's limb within the environment proves beneficial. Here, we want to better understand the relationship between haptic and visual feedback and how these two forms of sensory feedback interact during object acquisition.

Use of kinematic measures to infer planning

Movement time has been widely used to characterize the difficulty of a task in the area of HCI. This measure provides information regarding the difficulty of the movements, but does not give us a complete picture about the profile or shape of the movements being performed. In human movement studies, three-dimensional kinematic measures such as peak reaching velocity and deceleration time toward the target have long been used to characterize target acquisition movements [8]. Figure 1 illustrates a velocity profile for a typical reaching movement made to a physical target. Note the velocity profile of the movement resembles a bell shape: velocity increases to a single peak value and then decreases as the target is approached. These kinematic measures allow us to further understand how the movements are being planned and performed. As well they provide us with complementary measures of task precision.

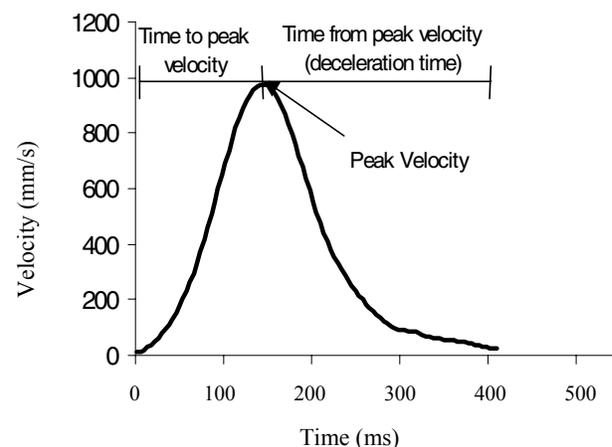


Figure 1: One trial showing a typical velocity profile for reaching movements. Note the asymmetrical bell shaped velocity profile.

MacKenzie, Marteniuk, Dugas and Eickmeier [7] performed a study replicating conditions of Fitts and Peterson's [2] discrete aiming movements. They replicated the systematic effects of target size on movement time described by Fitts & Peterson. However, MacKenzie et al. [7] also asked whether there was a reliable kinematic measure of the precision

requirements of the task. By differentiating the 3-D position data as shown in Figure 1, and then time normalizing the velocity profiles to 100 points for individual trials these authors discovered that as the diameter of the targets became smaller, the percent of time spent in the deceleration phase of the movement increased. MacKenzie et al. [7] have operationally defined this lengthening of the deceleration phase as the ‘precision effect’: as the precision required of the movement increases, deceleration time to the target increases as well. In the present experiment, we are also interested in using kinematic measures to further explore reaching movements in virtual environments. We expect that these measures will allow us to better understand how removing sensory information affects performance in a virtual environment.

This experiment was designed to address three purposes. First, we were interested in verifying that similar movement time results seen in typical aiming and docking experiments in computer generated environments would also be seen for *reaching to acquire* a computer generated target object. Second, we were interested in understanding how the availability of haptic and visual feedback affect movements made in augmented and virtual environments. Our third purpose was to use kinematic variables to obtain a more detailed understanding of how reaching movements are made in computer generated environments [3]. Our main research hypothesis was that haptic feedback at object contact would provide movement time and deceleration time benefits when acquiring a target into grasp. We also expected that the availability of visual and haptic feedback would interact such that subjects would have the slowest reaching speed when acquiring a graphic object without visual feedback of the moving limb. Finally, we expected that movement times would follow Fitts’ law for the various target sizes regardless of whether haptic feedback was available or not.

METHOD

Subjects

Eight university students were each paid \$10 for participating in a single, one-hour experimental session. All subjects were right-handed, and had normal or corrected-to-normal vision. Subjects provided informed consent. Ethical approval was obtained from the Simon Fraser University Ethics Committee.

Experimental apparatus

This experiment was conducted in the Enhanced Virtual Hand Laboratory (EVHL) at Simon Fraser University. Shown in Figure 2, the graphic image of a target cube produced by a Silicon Graphics Inc. (SGI) ONYX2 was displayed on a downward facing SGI RGB monitor. A half-silvered mirror was placed parallel to the computer screen, midway between the screen and the table surface. Thus, the image on the screen was reflected in the mirror and appeared to the subjects to be located in a workspace on the table surface.

The images for the left and right eye were alternately displayed on the SGI monitor and were synchronized with the CrystalEYES™ goggles worn by the subject. The subject thus obtained a stereoscopic view of the images being projected onto the mirror. Three infrared emitting diodes (IREDs) were fixed to the side frame of the goggles. A two-sensor OPTOTRAK 3020 motion analysis system (Northern Digital, Inc.) tracked the three dimensional position of the IREDs on the goggles at a sampling rate of 240 Hz. This information was processed by the SGI ONYX2, with a 20-40 ms lag, to provide the subject with a real time, head-coupled view of the image [19]. Finally, three IREDs were positioned on the subject’s thumb, index finger and wrist such the 3-D position coordinates of the movement of these landmarks could be tracked and stored for later analysis.

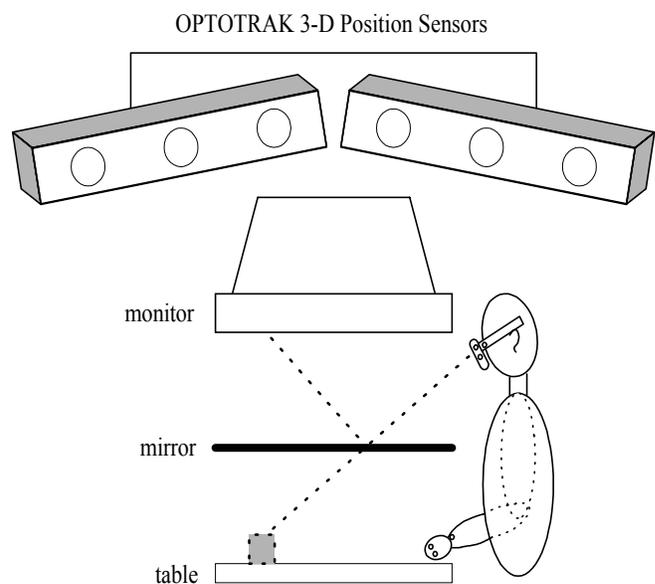


Figure 2: Illustration of the Enhanced Virtual Hand Laboratory. In this experiment, the target cube shown in grey could either be augmented (graphic and physical) or graphic only.

The experiment was conducted in a dark room. A light was positioned under the mirror to control visual feedback to the subject. When the light was on, the subject could see through the mirror, providing visual feedback of the moving limb and workspace below the mirror. When the light was off, the subject could see neither the workspace below the mirror nor the movement of the limb. In both conditions, a graphic representation of the target object was always available.

The target objects were shaded graphical cubes of four different sizes (12.7, 25.4, 38.1, and 50.8 mm³) located 15 cm directly in front of the start position of the subject’s hand. To ensure a comfortable grasping angle, the targets were also

rotated 45° clockwise about the vertical axis. In half the conditions, the graphic target object was superimposed over a physical wooden cube of the same size. In the other conditions, the graphic target only was presented so that the subject did not receive haptic feedback at object contact.

Experimental design

In the current experiment we manipulated target type, visual condition, and target size. In half the conditions, subjects reached to contact augmented cubes (physical cubes with superimposed graphic) while in the other conditions subjects reached for graphic cubes (no physical cube). The two visual conditions included the presence or absence of visual feedback of the limb and workspace below the mirror. With visual feedback, subjects had visual information about the movement of their limb, graphic information about the location of the target and visual information about the workspace below the mirror. Where visual feedback was absent, subjects had only graphic information about the size and location of the target. The workspace below the mirror was completely blacked out such that they were unable to see their moving limb. Thus, proprioception through muscle and joint receptors was the only feedback source available; proprioceptive feedback had to be integrated with vision to signal target acquisition. Finally, subjects reached to contact cubes of four different sizes. These manipulations resulted in a balanced design of 2 targets \times 2 visual conditions \times 4 cube sizes. Trials were counterbalanced across subjects on the visual condition, and target type; target size was randomized over trials. Six trials for each target size were presented in each experimental condition.

Procedure

At the beginning of the experiment, the subject was seated in a height-adjustable chair, in front of the tabletop virtual environment such that the forearm was at approximately the same height as the table surface. The subject was then asked to put on the CrystalEYES goggles. Individual subject's eye positions were calibrated relative to the IREDs on the goggles to give the subject a customized, stereoscopic view of the virtual environment. Deliberate steps were taken to ensure that the graphical target was accurately superimposed over the physical target for each individual in the augmented target condition. Subjects were asked to move the physical object such that it was superimposed over the graphical target. The chosen position for the physical object was recorded for each target size, and used in the remaining augmented trials to accurately position the physical target.

Subjects began each trial with the pads of the index finger and thumb lightly touching over a start position, and the remaining digits curled towards the palm. The task to be performed in every trial was to reach toward and grasp (but not lift) the target objects. Subjects were instructed to begin the movement when the graphical target appeared and to say "Okay" when the movement was complete.

Data analysis

OPTOTRAK 3-D position data from the wrist IRED were analyzed for specific kinematic measures. We were interested in measuring the following dependent measures: Movement Time (MT), Peak Velocity of the Wrist (PV), Time to Peak Velocity of the Wrist (TPV) and Percent Time from Peak Velocity of the Wrist (%TFPV). As discussed, movement time and percent time from peak velocity have typically been used to quantify task difficulty [2,7,10] while peak wrist velocity and the timing of that peak give us an indication of the shape of the movement [7,14].

Before extracting the dependent measures, the position data were interpolated, rotated into a meaningful coordinate system (x = forward movements, y = side to side movements, z = up and down movements) and smoothed with a 7 Hz low-pass second-order bi-directional Butterworth filter. A customized computer program was used to determine the start of movement based on a criterion velocity of 5mm/s [3]. The end of movement was determined as the point when the distance between the index finger and thumb IREDs did not change by greater than 2 mm over a period of 12 frames (50 ms). This stabilization of the distance between the fingers signified that subjects had completed their grasp. The position data were differentiated using customized software that performed a 5 point central finite difference technique. Peak resultant velocity and the timing of the peak were extracted using customized software. Percent time from peak velocity was defined as (MT-TPV)/MT*100. Data were analyzed using separate repeated measures ANOVAs and an a priori alpha level was set at $p < 0.05$. Means and standard error measures are reported for significant results.

RESULTS

Movement Time

For the time it took subjects to reach from the start position to make contact with the object, main effects were found for target type ($F_{1,7} = 20, p < 0.005$), visual condition ($F_{1,7} = 25.6, p < 0.005$) and cube size ($F_{1,7} = 4.1, p < 0.05$). Subjects took significantly longer to reach for a graphic cube (375 ± 12 ms) than an augmented cube (254 ± 8 ms). They also took longer to reach for an object when vision of the hand and workspace was removed (352 ± 13 ms) than when they were afforded full vision of the limb and workspace (277 ± 10 ms). Furthermore, as predicted by Fitts' law, subjects took longer to reach and grasp smaller targets than larger targets (small to large: $331 \pm 17, 311 \pm 17, 309 \pm 18, 307 \pm 18$ ms).

However, the main effects of target type and cube size have to be interpreted in light of a significant interaction of these two variables ($F_{3,21} = 8, p < 0.005$). As shown in Figure 3 movement times decreased as cube size increased, only in the augmented target condition. When the target was graphic only, movement time was similar across all four target sizes. Furthermore, the three-way interaction of target type \times visual condition \times target size did not reach significance levels ($F =$

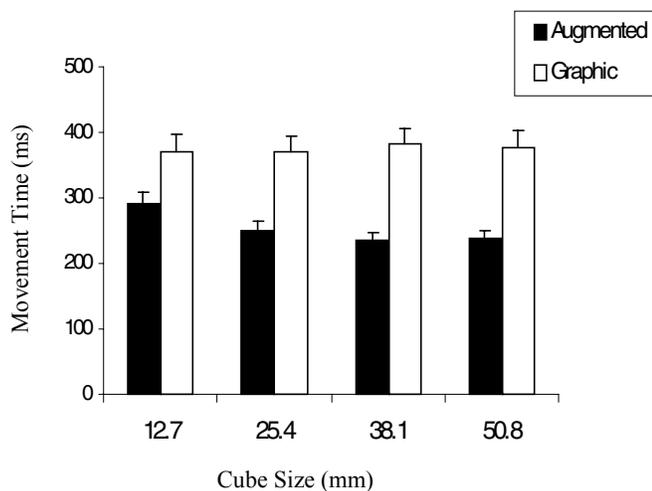


Figure 3. Interaction between target type and cube size on movement time.

1.2, $p > 0.1$) indicating consistent results whether subjects had visual feedback or not.

To assess whether our results support the notion that Fitts' law is followed in a *grasping* task, regression analyses on MT for both the augmented and graphic conditions using Fitts' original formulation were performed:

$$MT = a + b \log_2(2A/W), \text{ where } \log_2(2A/W) = ID$$

Results revealed a significant regression ($F_{1,62} = 8.5$, $p < 0.01$) for the augmented target condition, although a mediocre $r=0.35$ was found. The resulting regression equation was calculated to be:

$$MT = 157.4 + 94.2 ID$$

The low correlation value is probably due to the small number of indices of difficulty studied in this experiment as well as the proximity of the two smallest target IDs (ID = 2.56, 2.98, 3.56, 4.56). However, the significant regression is taken here as preliminary evidence that Fitts' law is followed when haptic feedback is available.

Results of the regression analysis for the graphic target condition failed to reach significance ($F_{1,62} = 0.08$, $p > 0.5$, $r=0.035$). Thus, at this time we have no evidence that Fitts' law holds for grasping tasks when haptic feedback is not available.

Peak Velocity

Velocity profiles, and specifically, the peak velocity attained when reaching give an indication of the shape of the movement. Consistent with the MT results, there were main effects of target type ($F_{1,7} = 6.9$, $p < 0.05$), visual condition ($F_{1,7} = 68.9$, $p < 0.001$) and cube size ($F_{1,7} = 8.9$, $p < 0.005$) on peak velocity. Figure 4 shows typical velocity profiles for the two target types and visual conditions for the smallest and largest targets. Note that peak velocity is higher for grasping an augmented target (mean = 1129 ± 34 mm/s) than

a graphic target (mean = 952 ± 28 mm/s) and for grasping under full vision (mean = 1160 ± 28 mm/s) than when vision is removed (mean = 920 ± 31 mm/s). Average peak velocities for the smallest to largest targets were 1096 ± 44 , 1067 ± 51.2 , 1019 ± 43 and 979 ± 29 mm/s respectively. Thus, as target size increased, peak velocity decreased slightly. Also note in Figure 4B that when reaching toward a graphic target without vision of the hand, the velocity profile was multi-peaked. Analysis of the percentage of trials during which multi-peaked velocity profiles were observed revealed that there was an interaction between target type and visual condition ($F_{1,7} = 9.70$, $p < 0.05$). Subjects produced significantly more multi-peaked velocity profiles when the lights were off and the target was graphic (43%) than for the other three conditions (Augmented/Lights On: 10%, Augmented/ Lights Off: 14.7%, Graphic/Lights On: 14.4%

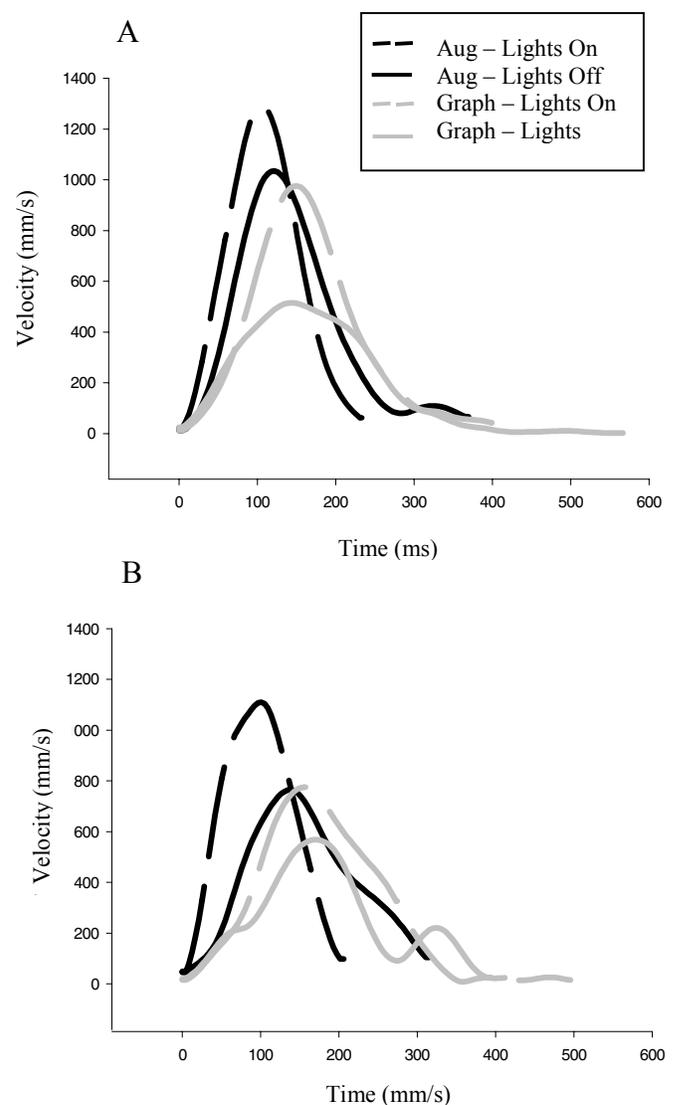


Figure 4. Typical velocity profiles for reaching movements made toward the smallest (A) and largest (B) targets.

trials with multi-peaked velocity profiles). This re-acceleration at the end of the movement might indicate movement corrections made to adjust for undershooting the target.

Time to Peak Velocity

A main effect for the timing of peak velocity was found for visual condition ($F_{1,7} = 47.6, p < 0.001$) which indicated that subjects reached peak velocity sooner when the lights were on (110 ± 2 ms) than when the lights were off (128 ± 3 ms). As well, an interaction between visual condition and object size was found ($F_{3,21} = 3.9, p < 0.05$). Figure 5 illustrates this interaction. Note that when the lights were off, the trend for time to peak velocity was to increase with object size, however, when the lights were on, the opposite effect was found.

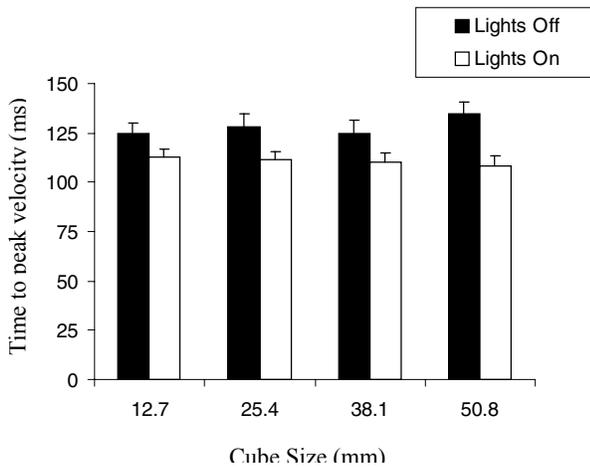


Figure 5. Interaction between target visual condition and target size on time to peak velocity.

Percent Time From Peak Velocity

For percent time from peak velocity, significant main effects were found for target type, visual condition and cube size. As well, each of these factors interacted with each other to

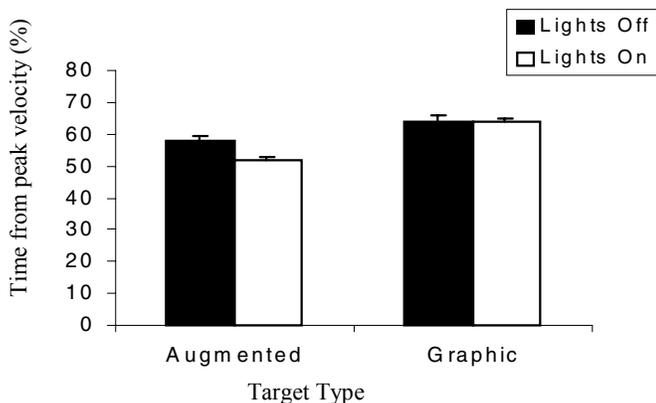


Figure 6. Interaction between target type and visual condition on percent time from peak velocity.

result in three two-way interactions: target type x visual condition ($F_{1,7} = 6.4, p < 0.05$), target type x cube size ($F_{3,21} = 3.9, p < 0.05$), visual condition x cube size ($F_{3,21} = 4.8, p < 0.05$). For brevity, only the three two-way interactions are discussed here. Deceleration time was always longer for reaching to a graphic target than an augmented target. However, when reaching to an augmented target, deceleration time was longer when the lights were off than when the lights were on. In contrast, when reaching to a graphic target, deceleration time was similar regardless of the presence or absence of visual feedback (see Figure 6). When reaching to grasp augmented targets of increasing size,

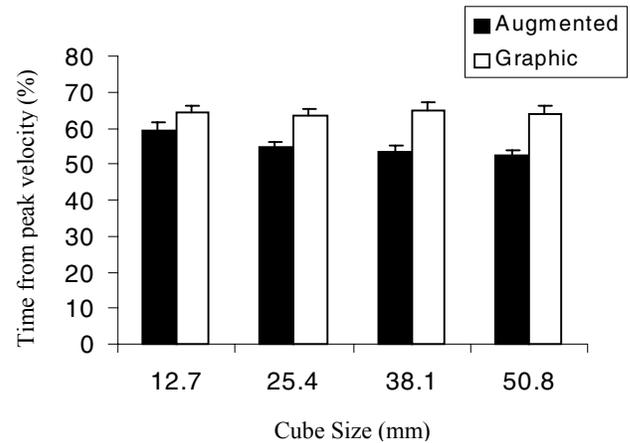


Figure 7. Interaction between target size and target type on percent time from peak velocity

deceleration time decreased. On the other hand, when reaching to grasp graphic targets, subjects had similar deceleration times regardless of cube size (see Figure 7). Figure 8 shows that deceleration time was always longer when visual feedback was not available. However in the absence of visual feedback, deceleration time decreased as target size increased. When visual feedback was available, deceleration time was similar regardless of target size.

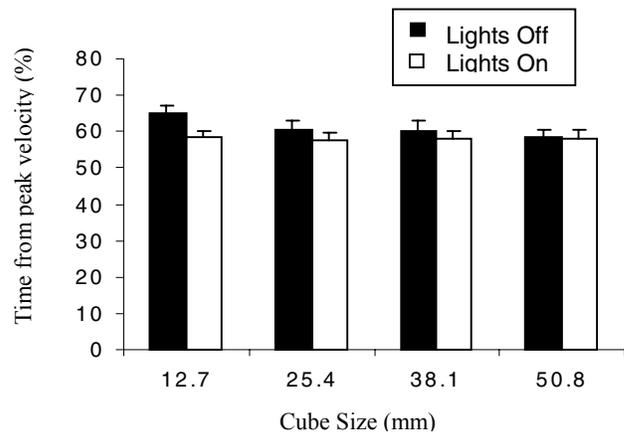


Figure 8. Interaction between target size and visual condition for percent time on peak velocity.

DISCUSSION

In this experiment, we studied how the availability of haptic and visual feedback affect reaching movements to acquire an object in a desktop virtual environment. We have shown that both haptic and visual feedback have profound effects on human performance.

In the past, Fitts' law has been shown to hold under a variety of conditions. Studies have been conducted that have replicated Fitts' law in space, underwater, and even with limbs other than the hand [see 17]. Thus, Fitts' law has been found to be quite robust under most circumstances. However, in the present study, we have shown that Fitts' law does not always hold when making reaching movements towards objects in virtual environments. Our results indicate that when subjects reached to grasp objects that had only a graphic representation, movement time was the same regardless of object size. These results were found whether subjects had visual feedback of the moving limb or not.

Why did Fitts' law not hold when haptic feedback was removed? This result is contrary to our hypothesis and indeed quite puzzling. In a study conducted using real targets in a natural environment, MacKenzie [9] replicated Fitts' law regardless of whether a haptic stop was available to indicate target acquisition or not. As well, Wang, et al. [20], have shown that in virtual environments, Fitts' law holds regardless of the presence of haptic constraints. One major difference between these two studies and the current experiment, is in the goal of the task. In MacKenzie [9], the task goal was to aim to a target, and in Wang et al. [20] the task goal was to dock an object in hand. But, in the current experiment, subjects *reached to acquire* an object into grasp.

It has been shown that task goal does indeed influence aiming movements in natural environments [14][15], and the results shown here seem to indicate the same result for computer generated environments. Perhaps because of the terminal accuracy required to accomplish the task in this experiment, haptic feedback became an essential source of information about task completion. Thus, when haptic feedback was not available, a ceiling effect occurred, and subjects took longer regardless of object size to acquire the target. Further research is needed to elucidate this important effect.

The role of visual feedback about the movement of the limb in the surrounding environment was also investigated in the current experiment. Consistent with the findings of Graham and MacKenzie [3] and Mandryk [12], movement time was significantly reduced when vision of the moving limb was permitted. As well, we saw that deceleration time was shortened when subjects were able to see their limbs move within the environment. These results indicate a need to provide users with some representation of their movements in order to improve performance.

Implications for HCI

Our results indicate that in order for humans to work effectively in virtual environments, some form of haptic and

visual feedback should be included in these systems. Recently, force feedback devices have been implemented in virtual environments to enhance the realism of interaction [13]. While it is believed that the addition of haptic feedback improves performance in virtual environments, there has been little empirical evidence to support this claim [1]. The results from the current experiment lend further empirical support to the notion that haptic feedback, especially with respect to object contact is crucial for humans to produce optimal performance in computer generated environments. Our results show performance improvements in terms of reduced movement time when haptic and visual feedback are available. They also do not provide evidence that a fundamental law of human movement, specifically Fitts' law holds when haptic feedback is unavailable in object acquisition tasks. These two results confirm that we must pay more attention to the use of sensory feedback in virtual environments in order to capitalize on the human ability to manipulate physical objects with ease.

Use of kinematic variables has also provided us with a powerful tool to study how movements are made under various conditions. By looking at the velocity profiles, we were able to determine that in simple conditions, movement profiles in computer generated environments resemble the bell-shaped movements made to simple targets in natural environments. However in the more complex task of reaching without vision to a graphic target, we saw a multi-peaked velocity profile. This multi-peaked profile indicates that subjects made corrective movements toward the end of their reach. As well, by measuring the timing components of the movement, specifically time to peak velocity and percent time from peak velocity we were able to gather more information about the shape of the movements being made. Our results indicate that the shape of the movement, such as when peak velocity occurs and how much time is spent in deceleration, depends on the availability of haptic and visual feedback as well as the size of the target. This has serious implications for the design of augmented environments and for implementing movement prediction algorithms to improve the speed of graphics in interactive computer systems. By using data from human movement studies, we may be able to mathematically model and predict upcoming movements. Kinematic information about the shape of the movements will be *essential* to accomplish this goal.

In conclusion, we have shown that visual information about the movement of the limb in the environment and haptic feedback about object contact have critical effects on human performance in virtual environments. We recommend that in order to optimize human performance in computer generated environments, attention should be paid to providing the user with veridical haptic and graphic sensory information about their movements within the environment.

ACKNOWLEDGEMENTS

This research was supported by grants from the Natural Science and Engineering Research Council of Canada.

REFERENCES

1. Arsenault, R. and Ware, C. (2000). Eye-Hand coordination with force feedback. *Proceedings of the Conference on Human Factors in Computing Systems CHI '00*, ACM Press, 408-414.
2. Fitts, P.M., & Peterson, J.P. (1964). Information capacity of discrete motor response. *Journal of Experimental Psychology*, 67(2), 103-112.
3. Graham, E.D. and MacKenzie, C.L. (1996). Physical versus virtual pointing. *Proceedings of the Conference on Human Factors in Computing Systems CHI '96*, ACM Press, 292-299.
4. Hoffman, H.G. (1998). Physical touching of virtual objects using tactile augmentation enhances the realism of virtual environments. *Proceedings of IEE VRAIS*, 59-63.
5. Klatzky, R.L. & Lederman, S.J. (1999). The haptic glance: A route to rapid object identification and manipulation. In D. Gopher & A. Koriats (Eds.) *Attention and Performance XVII. Cognitive regulations of performance: Interaction of theory and application*, 165-196. Mahwah, NJ: Erlbaum.
6. Linderman, R.W., Sibert, J.L, and Hahn, J.K. (1999). Towards usable VR: An empirical study of user interfaces for immersive virtual environments. *Proceedings of the Conference on Human Factors in Computing Systems CHI '99*, 64-71. New York: ACM Press.
7. MacKenzie, C.L., Marteniuk, R.G., Dugas, C. and Eickmeier, B. (1987). Three-dimensional movement trajectories in Fitts' task: Implications for motor control. *The Quarterly Journal of Experimental Psychology*, 39A, 629-647.
8. MacKenzie, C.L. and Iberall, T. (1994). *The Grasping Hand*. Amsterdam: Elsevier Science.
9. MacKenzie, C.L. (1992). Making contact: Target surfaces and pointing implements for 3D kinematics of humans performing a Fitts' task. *Society for Neuroscience Abstracts*, 18, 515.
10. MacKenzie, I.S. (1995) Movement time prediction in human-computer interaction. In R.M. Baecker, J. Grudin, W.A.S. Buxton & S. Greenberg (eds.) *Readings in Human Computer Interaction: Toward the Year 2000*, 483-493, Morgan Kaufmann Publishers, San Francisco Ca
11. MacKenzie, I.S., Sellen, A. and Buxton, W. (1991). A comparison of input devices in elemental pointing and dragging tasks. *Proceedings of the Conference on Human Factors in Computing Systems CHI '91*, 161-166. New York: ACM Press
12. Mandryk, R.L. (2000). Using the finger for interaction in virtual environments. M.Sc. Thesis, School of Kinesiology, Simon Fraser University, Vancouver, B.C., Canada.
13. Mark, W.R., Randolph, S.C., Finch, M., Van Verth, J.M. and Taylor, R.M. (1996) Adding force feedback to graphics systems: issues and solutions. *Proceedings of the 23rd annual conference on Computer graphics*, 447-452.
14. Marteniuk, R.G., MacKenzie, C.L., Jeannerod, M., Athenes, S., and Dugas, C. (1987). Constraints on human arm movement trajectories. *Canadian Journal of Psychology*, 41, 365-378.
15. Mason, A.H., and MacKenzie, C.L. (2000). Collaborative Work Using Give-and-Take Passing Protocols. *Proceedings of the XIVth Triennial Congress of the International Ergonomics Association and the 44th Annual Meeting of the Human Factors and Ergonomics Society, Human Factors and Ergonomics Society*, 519-522.
16. Prablanc, C., Echaliier, J.F., Komilis, E. & Jeannerod, M. (1979). Optimal response of eye and hand motor systems in pointing at a visual target. I. Spatiotemporal characteristics of eye and hand movements and their relationships when varying the amount of visual information. *Biological Cybernetics*, 35, 113-124.
17. Rosenbaum, D.A. (1991) *Human Motor Control*, Academic Press, Toronto.
18. Sturman, D.J., Zelter, D. (1993). A design method for "Whole-Hand" human-computer interaction. *ACM Transactions on Information Systems*, 11,(3), 219-238.
19. Swindells, C., Dill, J.C., Booth, K.S., (2000). System lag tests for augmented and virtual environments. *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology: CHI Letters*, 2(2), 161-170.
20. Wang, Y. and MacKenzie, C.L. (2000). The role of contextual haptic and visual constraints on object manipulation in virtual environments. *Proceedings of the Conference on Human Factors in Computing Systems CHI 2000*, ACM Press, 532-538.