

Evaluating the Influence of Haptic Force-Feedback on 3D Selection Tasks using Natural Egocentric Gestures

Vijay M Pawar *

Department of Computer Science, University College London

Anthony Steed†

Department of Computer Science, University College London

ABSTRACT

Immersive Virtual Environments (IVEs) allow participants to interact with their 3D surroundings using natural hand gestures. Previous work shows that the addition of haptic feedback cues improves performance on certain 3D tasks. However, we believe this is not true for all situations. Depending on the difficulty of the task, we suggest that we should expect differences in the ballistic movement of our hands when presented with different types of haptic force-feedback conditions. We investigated how hard, soft and no haptic force-feedback responses, experienced when in contact with the surface of an object, affected user performance on a task involving selection of multiple targets. To do this, we implemented a natural egocentric selection interaction technique by integrating a two-handed large-scale force-feedback device in to a CAVETM-like IVE system. With this, we performed a user study where we show that participants perform selection tasks best when interacting with targets that exert soft haptic force-feedback cues. For targets that have hard and no force-feedback properties, we highlight certain associated hand movement that participants make under these conditions, that we hypothesise reduce their performance.

Keywords: Haptics, 3D selection, task performance, two-handed interaction, force-feedback

Index Terms: H.5.2 [Information Interfaces and Presentation]: User Interfaces—Haptics; I.3.6 [Computer Graphics]: Methodology and Techniques —Interaction techniques;

1 INTRODUCTION

One of the fundamental building blocks of any interactive system is the user interface for selecting objects. This is particularly true for virtual environments (VEs), where to manipulate a 3D object, we must first be able to select it. Selection, that is the specification of the object of interest, is a common precursor to subsequent tasks [19]. To date there has been considerable work on designing selection techniques suitable for VEs. Popular and commonly implemented examples include virtual hand, where the user touches the object of interest and ray casting where the user points at the object of interest (see [4] for a review).

From this large body of work, we can identify several factors that contribute to the performance of selection techniques: the devices being used, the hand being used and the type of setup [3] [12]. In contrast, there is little research evaluating the affects of additional sensory modalities. Of the few studies that do exist, the main focus of interest is the addition of haptic force-feedback cues, considered to be important for direct manipulation of objects when using our hands [2] [13] [21] [11].

It has been suggested that the extra bandwidth provided by haptic cues will benefit user performance. Starting from an information

theory perspective such as Fitt's Law [7], one can argue that the addition of haptics should provide a faster time to selection because the additional feedback gives more information to the user. Additionally, as the latency of haptic output is less than visual cues, users will receive information of their surroundings fractionally earlier, thus aiding their performance. However, we believe this benefit is not realised in all situations. From mechanical point of view, when confronted with 3D objects that exert haptic force-feedback, users must put in more physical work as they either have to push through the target object that provides resistance, or avoid objects thus taking a longer path to complete the task. In a single selection task, which forms the basis of most Fitt's-style studies, we might expect the former factor to dominate. However, we expect that the latter factors will come more relevant when selecting multiple targets.

We have integrated a two-handed large-scale force-feedback device in to a CAVETM-like IVE system. To ensure intuitive representation of the haptic cues, we co-located the haptic device to the tracking system, placing the 3D virtual haptic contact points at the tip and "underneath" the participants index fingers; thus implementing a virtual-hand like interaction technique [20]. We can provide two types of feedback: hard haptic feedback that prevents the user's finger(s) penetrating the target, and soft haptic feedback which provides resistance to the user's finger(s). These calibration and implementation details are minor contributions of the paper.

The main contribution of this paper is a study investigating how different types of force-feedback cues affect selection times as the difficulty of the selection task increases. Unlike previous studies that focus on measuring selection times only, we also evaluate how haptic force-feedback changes the user's strategy when completing the selection of multiple targets. From this, we identify how these hand movements affect the user, which we hypothesise may help and hinder their overall task performance.

2 RELATED WORK

Recently, with the availability of affordable haptic devices, there has been more research into 3D interaction techniques that use force-feedback cues. For VEs, work introducing haptics for selection metaphors have tended to use small-scale devices such as the PHANToM. Wall et al. investigated if the addition of haptic force-feedback, gravity wells and stereo graphics benefited the selection of objects in 3D [22]. They showed that accuracy improved, but not performance time, while stereo graphics helped both significantly. Magnusson et al. also used a PHANToM in a memory game showing that conditions with force-feedback had the best results [16].

In spite of the large number of human factor studies, only few measure the performance of a haptic interface or haptic feedback itself [18]. Wall and Harwin [23] employed a tapping test [14] in order to establish a measure of human performance in a simple target selection task. They showed that force-feedback significantly reduced subjects' movement times. In general, most studies compare the performance of haptic devices using tests such as 3D peg-in-holes and rendering of hard virtual surfaces [5].

At present, the rendering of hard virtual surfaces [10] has been the most common benchmark topic in evaluating the performance of haptic interfaces. Guerraz et al. [8] suggested to use physical data from a haptic device to evaluate haptic user interfaces. Kap-

*e-mail: v.pawar@cs.ucl.ac.uk

†e-mail: a.steed@cs.ucl.ac.uk

pers et al. [15] performed haptic identification experiments using quadric surfaces and showed that both shape index, a quantity describing the shape, and the curved nature of the surface had significant effect on haptic shape identification. Building upon these studies, Kirkpatrick and Douglas [9] produced similar work. They expressed their results in bits of information transfer and showed that humans could correctly identify at most 3 to 4 sphere sizes (corresponding to 2 bits) ranging from 10 to 80 mm in radius using the PHANToM. Murray et al. [29] also used an information transfer concept to evaluate their wearable vibrotactile glove.

Unlike previous work, we have implemented a technique that simulates natural pointing gestures for full scale arm movements. Due to the large movement range, the results on small-scale devices are not directly transferable, but given the information theoretic arguments presented above, we can expect haptic force-feedback cues to have a significant effect on our 3D selection tasks.

3 HYPOTHESIS

Our first hypothesis is that, as others have found when using small-scale haptic force-feedback devices, that the combination of haptic and visual feedback will be superior to visual feedback on single selection task performance (time to touch a single object from a starting position). The second hypothesis is that this will no longer be true for complex selection tasks because the haptic force-feedback will cause the user to take longer or slower paths to targets after selecting the first. A more formative hypothesis is that we expect the experience of soft or hard haptic force-feedback cues to have different effects on the user, but we do not have a prior expectation of which will produce faster performances.

Ideally within IVEs we wish to always support unbiased two-handed selection. However, in many facilities there is often a constraint, such as only one hand-tracker or one force-feedback device being available. Thus we will run both single (right), and two-handed trials. To scope the work, we are will only study virtual hand techniques, where the user touches the object they wish to select.

4 METHODOLOGY

4.1 Apparatus

To test our hypothesis, we created an IVE containing both visual and haptic content. For synchronised rendering of these input cues, we used two separate systems to individually maintain the high refresh rates needed: a CAVETM-like projection display (see section 4.1.1) and a GRAB haptic interface (see section 4.1.2). When combining these two systems together to display one continuous VE, we registered their local co-ordinate frames into a common spatial and temporal domain; co-locating the 3D haptic contact points (the virtual points where force-feedback is felt when using the GRAB haptic interface) precisely to the tips of both our index fingers (see section 4.1.3). By doing so, we ensured that participants could interpret and interact with the displayed content in an intuitive manner. Additionally, we chose hardware with large working volumes to allow the use of both hands, and ballistic movements over all distances within arm's reach. With this setup, we presented an egocentric interaction technique very similar to real world pointing and touching of objects. Overall, by changing the types of force-feedback perceived through the haptic arms (also see section 4.1.3), we evaluated the effect of different haptic force-feedback conditions on our ability to perform 3D selection tasks whilst using this "natural" selection interaction technique implemented.

4.1.1 Rendering Visual Sensory Input Cues

We rendered the visual sensory input cues using the ReaCToR at University College London, which is a CAVETM-like [6] display system. This consisted of three 3m x 2.2m walls and a 3m x 3m




Figure 1: Typical configuration of the ReaCToR with GRAB haptic interface. BOTTOM: GRAB haptic interface with VICON markers

floor, stereo projected from the back and top respectively. This is shown in Figure 1. The visuals were generated by self-built PC cluster consisting of a cluster master node with 2GB RAM and dual 1.8GHz Intel processors, and four cluster slave nodes with 1GB RAM, single 2.7GHz Intel processors and GeForce Quadro 5600 graphics cards. All cluster nodes run Windows XP. Each participant wore CrystalEyes stereo glasses, which we tracked using an Intersense IS900 system. With this, we could dynamically change the viewpoint of the virtual environment in relation to the participant's position and head rotation by performing low-latency precision head-tracking, allowing distortion-free and lag-free movement calibrated to the dimensions of the surrounding walls. The ReaCToR runs at a maximum refresh rate of 85Hz, stereo at 42.5Hz. The end to end latency is approximately 80ms.

4.1.2 Rendering of Haptic Sensory Input Cues

To simulate haptic cues, we used a GRAB haptic interface [1] as shown in Figure 1. This is a force-feedback device designed for the simulation of manipulation of objects with two fingers. This system consisted of two robotic arms and a control unit connected to a visualisation system (not shown in Figure 1).

We placed the GRAB haptic interface in the centre of the ReaCToR, 1m apart. To use the device, the participants placed their index fingers in the two thimbles placed at the end of each arm. This offered three degrees of freedom for both hands. The arms also tracked the position of each hand and sent this data to a visualisation package for representation within its scene (Figure 2). Additionally, we placed VICON markers on each arm, which we used to align and co-locate the virtual haptic contact point to the physical position of the thimble (see section 4.1.3). Overall, the workspaces of the two devices overlapped, so the participant could move and touch their hands together in a subset of the full working volume in front of them. Set at 600x400x400mm, this was large enough to provide haptic force-feedback over almost the full range of arm motion in front of the participant.

To exert force-feedback over the three degrees of freedom offered, each robotic arm is actuated using three DC motors and tendon transmissions. This gave a large peak and continuous force of 10-15N and 4-6N respectively. The device had a high stiffness of 3-7N/mm and a low mechanical friction of 20mN and 200mN (with active weight compensation). With this, a high degree of isotropy is

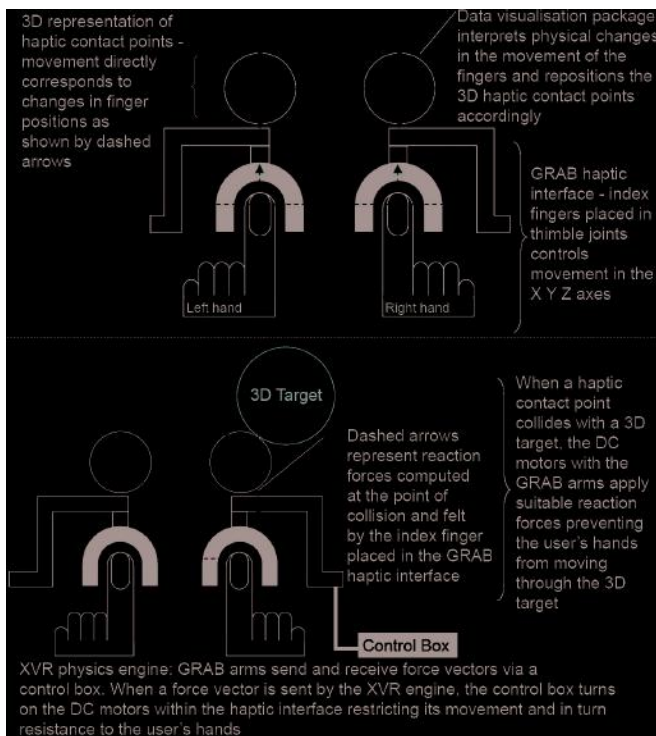


Figure 2: Illustration of co-located 3D haptic contact point to thimble joint. BOTTOM: Application of forces with GRAB haptic interface

present allowing a uniform reflected inertia of 0.4Kg throughout the workspace. We controlled the forces applied to each arm through the data visualisation package. This sent force vectors to the haptic control box that in turn switched on the DC motors directing the movement of the arms.

For this paper, we simulated two types of force behaviours: objects with hard surfaces and objects with soft surfaces. By using a physics engine within the visualisation package we computed the necessary reaction forces in real-time for the arms to simulate these two types of behaviours appropriately. For the simulation of hard surfaces, we computed a large reaction force of 5N the instant a 3D haptic contact point intersected the surface of a target. To simulate a soft surface behaviour, the reaction force sent to the haptic control box started at a small value gradually increasing to 3N using a linear function as the haptic contact point got closer to the centre of the target. We piloted these behaviours to determine the correct reaction force values specifically for this device to give reliable and informative feedback. These force values may not be suitable for other devices.

4.1.3 Integration of Visual and Haptic Sensory Input Cues

To render both visual and haptic content within the VE, we used the XVR visualisation software [17]. We configured this package specifically for the dimensions of the ReaCToR, head tracking system and GRAB haptic arms used in this study.

For participants to interpret the inputs cues generated by each haptic arm in a coherent manner with the visuals, we aligned the devices' local co-ordinate frames to the reference frame of the XVR software. To do this, we placed VICON markers around the thimble joint and base of each haptic arm as shown in Figure 1. By using an array of infra-red VICON cameras placed at ceiling level around the ReaCToR, the VICON IQ software and its real-time engine computed the 3D positions of these markers in relation to the dimensions of the surrounding walls. From this information, we




Figure 3: Examples of the selection task being performed. LEFT: Participant within the ReaCToR. RIGHT: An example of the configuration of the world in a selection trial

repositioned the 3D haptic contact points precisely to the tip of the thimble joint for each arm (Figure 2). We found a position of 1cm in front of the hands preferable as it avoided occlusion errors during movement, yet was close enough to feel intuitive. Additionally, to account for image distortions associated with instability with the stereo and head tracking as participants rotate and move their head, we set the shape and size of the 3D haptic contact points to a small grey sphere of 2.5cm in radius for each arm.

To define the behaviour of the 3D haptic contact points, we also used the XVR visualisation package to control their movement. As we wanted to display a natural interaction technique we mapped the transformation of these grey spheres directly to the physical movements of the haptic arms in all axes in a 1-1 manner. Similarly, we used the XVR physics engine to compute suitable reaction forces when defining soft and hard force-feedback responses for each haptic arm to exert as shown in Figure 2. Again, as we wanted to focus on natural interaction techniques, we used a real world model to define the forces and behaviour of the physics within the VE. As a result, each arm applied forces onto the participant only when their associated 3D haptic contact point would intersect the surface of a surrounding 3D object.

4.2 Experimental Design

With the above setup, we presented an experiment where we asked participants to perform a series of 3D selection tasks. By using the haptic arms and the egocentric pointing technique offered, participants could manoeuvre the finger tips of either one or both their hands, and select a series of targets placed within arm's reach (Figure 3). To evaluate the influence of haptic force-feedback when performing these tasks, we tested three different force-feedback conditions whilst keeping the visual input cues constant:

1. No Haptic Force-Feedback- No force-feedback cues when in contact/selecting a target.
2. Hard Haptic Force-Feedback- A hard force-feedback response when a target is selected, similar to touching a wooden or marble table.
3. Soft Haptic Force-Feedback- A soft force-feedback response when in collision with a target, similar to pressing on a cushion or sponge.

The experiment consisted of one blue, one red and one yellow coloured sphere targets placed on grey rods within an environment. As shown in Figure 5, we placed these targets within an outdoors scene as we wanted to infer real-world responses whilst interacting within the presented VE. Additionally, this gave a fixed horizon level which helped reduce any adverse side effects caused by simulator sickness. By using this setup, we positioned these sphere targets randomly in front and always within arm's reach of the participant to select. For each force-feedback condition tested, we asked every participant to perform a series of selection tasks covering 3 difficulty classes:

Hands Used	Number of selection trials for task class								
	No Feedback			Soft Feedback			Hard Feedback		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
Right	15	15	15	15	15	15	15	15	15
Both	15	15	15	15	15	15	15	15	15

Table 1: Overview of experiment. S1-selection of 1 target type of task, etc.

1. Selection of one target ('Select 1')- Only one object in the scene - one blue sphere. Participants were asked to select the blue sphere.
2. Selection of two targets ('Select 2') - Only two objects in the scene - one blue and one red sphere. Participants were asked to select the blue sphere and then the red sphere.
3. Selection of three targets ('Select 3')- Three objects in the scene - one blue, one red and one yellow sphere. Participants were asked to select the blue sphere, then the red sphere and finally the yellow sphere.

For each of the above groups, we identified 15 predefined random sphere positions distributed uniformly within the workspace. We displayed each of these individual selection tasks from each group all together in a random order. To further avoid any outside effects on the interaction of the task, we did not use any distracting factors [22], and set the size of the spheres to 10cm for all targets. Additionally, depending on the haptic condition tested, all targets had the same physical properties.

In the two-handed tasks the participants could use either hand to select the next target. In the right-handed tasks, they needed to reach all the targets with that hand. When selected, the target would turn grey- a common visual selection cue. Once all the one, two or three targets were selected in the correct order, two reset markers would appear for both hands to touch and reposition their hands to where they started. At this point, we would automatically load a new selection task to perform. This process repeated until participants completed all tasks covering all three selection task classes.

We tested two types of hand interactions: performing selection tasks with both hands and with only the right hand. Table 1 gives an overview of the conditions we assessed for each of the hand interaction types tested. Each individual participant used only one force-feedback condition, but did this with both-handed and right-handed interaction types. They first performed 45 selection trials with only their right hand, and then did another 45 trials using two hands. There was a 15 minute break between these two sets of 45 trials. In total, each participant performed 90 selection tasks.

4.2.1 Measurements

No guidelines or benchmark tests yet exist for measuring haptic force-feedback so we took a thorough approach to recording the performance of each participant tested. In addition to taking classical user performance data such as time taken, velocities and distance travelled to complete each selection task, we also asked participants to fill out usability and presence questionnaires [11].

5 RESULTS

To evaluate the effects of the three haptic force-feedback conditions tested, we analysed the user performance when selecting 1, 2 and 3 targets by using quantitative and qualitative data. In the tables and figures presented, we use the following terms and labellings: 'Select 1'-selection of 1 target class, 'Select 2'-selection of 2 targets class, 'Select 3'-selection of 3 targets class. Figures 4 and 5, 'Blue bar'-time or distance travelled to select blue target, 'Red bar'-time or distance travelled to select red target, 'Yellow bar'-time or distance travelled to select yellow target. Figure 6, 'Blue Line'-hard force-feedback cues, 'Green Line'-soft force-feedback cues, 'Red Line'-no force-feedback cues. Term 'FB'-force-feedback.

Right-Handed Interaction			
Time (s)			
	Select 1	Select 2	Select 3
Hard FB	6.70	12.74	21.02
Soft FB	6.03	11.44	17.23
No FB	7.02	12.40	18.82

Two-Handed Interaction			
Time (s)			
	Select 1	Select 2	Select 3
Hard FB	7.14	12.74	16.83
Soft FB	6.25	11.44	14.12
No FB	7.64	12.40	17.87

Table 2: Average time taken for each class of selection task

5.1 Participants

In total we had 45 participants (33 male and 12 female). From the questionnaires completed, all participants were of similar age (20-25) and backgrounds. All were right-handed with good hand-eye coordination, and all had a good appreciation of 3D games. Before starting the experiment we gave each a demonstration of the equipment and thorough instructions. Each participant had 10-15 minutes to accustom themselves with the GRAB haptic interface, ReaCTor, head tracking and egocentric pointing interaction technique to level out any learning effects. Once done, we repeated the instructions, answered any questions, and asked if they were ready to continue with the experiment. We logged measurements during the experiment, and when finished asked each participant to complete usability and presence questionnaires. 15 participants completed each of the three haptic force-feedback conditions. Recall that each subject did 45 right-handed selection trials first and then 45 two-handed selection trials with a 15 minute break in between. The results are thus presented in as between subjects comparisons of the force-feedback methods, or separated on a particular class of selection types (Select 1, Select 2, Select 3).

5.2 Quantitative Data

5.2.1 Average Time Taken

For each of the three classes of selection tasks, we computed the average time it took for participants to select the presented targets. As shown in Figure 4 and summarised in Table 2, we can see the quickest performances for all selection tasks when using either right or two-handed interaction came under the soft force-feedback conditions. To evaluate this further, we calculated the differences in time between each force-feedback condition, presented Table 3. This showed, for both right and two-handed interaction, firstly, that a soft force-feedback condition produced the best results: on average 1.3 and 2.7 second faster than hard force-feedback conditions for tasks involving the selection of 2 and 3 targets respectively. Compared with no force-feedback conditions, soft force-feedback was quicker by over 1 second on average but no more than 3 seconds for all three classes of selection tasks.

With respect to trends between hard and no force-feedback conditions this was varied. Again, by looking at the differences in time taken when selecting 2 and 3 targets as shown in Table 3, we can see that participants performed marginally better under hard force-feedback conditions when using two hands: on average 0.83 and 0.80 seconds faster. Conversely this was not the case when selecting targets with the right hand only. On average, a hard force-feedback condition produced slower times: by 0.34 and 2.08 seconds when selecting 2 and 3 targets respectively in comparison to a no force-feedback condition.

To assess the significance of results in Figure 4 and Tables 2 and 3, we performed a single factor ANOVA for each of the 45 tasks tested. We performed separate ANOVAs because the individual tri-

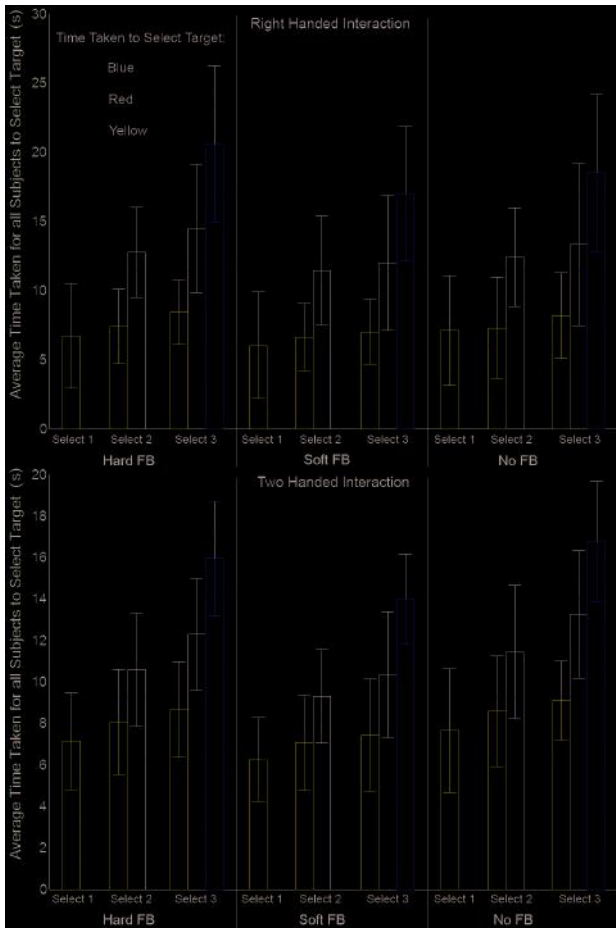


Figure 4: Average time taken for each selection task class. See note at beginning of Section 4 for an explanation of the labelling

als vary quite significantly in complexity. In table 4 we present a set of comparisons showing the number of tasks out of 15 that have p-values less than 0.05 between each haptic condition tested.

By combining the data in Tables 2, 3 and 4, we see that participants take less time to select targets under soft force-feedback conditions. In comparison to data collected under hard force-feedback conditions, this is particularly evident, as the number of occurrences when the results are significantly faster increases from 5 tasks to 11 and 1 to 11, for right and two-handed interaction respectively as the difficulty of the task increases over all three selection task classes. With respect to no force-feedback, the number of tasks indicating significantly better results under soft conditions appears in 4 tasks on average for all three classes of selection tasks. Again this shows that participants achieved faster times in the soft feedback condition, but not as frequently and it does not increase with the difficulty of the task. Consequently, this suggests that these results between soft and no force-feedback overlap. By looking at the standard deviation bars in Figure 4, we can see this is the case at least for the selection of 1 and 2 targets.

With respect to hard force-feedback conditions versus no force-feedback conditions, participants performed differently depending on whether they used two hands or just the right hand to select the presented targets. As shown in Table 4, when the difficulty of the task increases, so does the number of tasks where the results are significantly different: from 0 to 8 for two-handed interaction and 1 to 7 for right-handed interaction. By looking into this further, and comparing the data in Table 3, participants performed the se-

Right-Handed Interaction			
Time (s)			
	Select 1	Select 2	Select 3
Soft FB vs Hard FB	-0.67	-1.31	-3.57
Soft FB vs No FB	-1.05	-0.96	-1.49
Hard FB vs No FB	-0.38	0.34	2.08

Two-Handed Interaction			
Time (s)			
	Select 1	Select 2	Select 3
Soft FB vs Hard FB	-0.88	-1.30	-1.97
Soft FB vs No FB	-1.38	-2.13	-2.77
Hard FB vs No FB	-0.50	-0.83	-0.80

Table 3: Average difference in time taken between each force-feedback condition tested

Right-Handed Interaction			
	Select 1	Select 2	Select 3
Hard FB vs Soft FB	5	8	11
Hard FB vs No FB	0	1	7
No FB vs Soft FB	9	3	4

Two-Handed Interaction			
	Select 1	Select 2	Select 3
Hard FB vs Soft FB	1	6	11
Hard FB vs No FB	0	1	8
No FB vs Soft FB	4	3	4

Table 4: Pair-wise comparison of time taken data: number of tasks out of 15 within selection class with significantly different results

lection of 3 targets 0.8 seconds faster with two hands with hard force-feedback, whereas, with only the right hand it is 2.08 seconds slower. Considering this difference is large, we suspect there are other contributing factors to this result such as how participants are moving their hands under these conditions.

5.2.2 Average Distance Travelled

For each of the three classes of selection tasks, we computed the average distance travelled for participants to select the presented targets. As shown in Figure 5, we can see the shortest distances achieved for all selection tasks when using either right-handed or two-handed interaction came under soft and hard force-feedback conditions. Again, evaluating this further, we calculated the differences in distance between each force-feedback condition, as shown in Table 6. This shows that participants under soft and hard force-feedback conditions for both types of hand interactions produced trajectories that took on average 0.2 and 0.30 metres less distance in comparison to no force-feedback conditions for tasks involving the selection of 2 and 3 targets respectively.

As we can see from Figure 5 and summarised for total distance in Table 5, for all selection tasks performed under soft and hard force-feedback, the distance travelled is much smaller compared to the condition where there is no force-feedback. From table 6, we can see that as the difficulty of the selection task increases to 3 targets, so does the difference travelled by participants under soft force-feedback conditions in comparison to hard force-feedback conditions: 0.42m and 0.39m less for two-handed interaction and 0.27m and 0.26m for right-handed interaction respectively. Again, to assess the significance of these trends, we performed a single factor ANOVA for each of the 45 tasks tested. We presented a set of comparison tables showing the number of tasks that have p-values less than 0.05 between each haptic force-feedback conditions tested.

By comparing data from Tables 5,6 and 7, this confirms the trend that participants take a longer path when selecting objects without haptic force-feedback. As the difficulty of the selection task increases from 2 to 3 targets, all 15 tasks showed significant differences between conditions with and without haptic force-feedback




Figure 5: Average distance travelled for each selection task class

for both types of hand interaction. With respect to differences between soft and hard conditions, the data showed at times a soft condition produced better results, however this was not evident for all selection tasks. At best, the differences between soft and hard conditions were at most 0.06m, which is small. Overall, from this data, it suggests participants move their hands differently when haptic force-feedback is present compared to when it is not, especially for selection tasks involving more than 1 target.

5.2.3 Velocity

To link the results for time and distance taken together, we computed the average velocity for each of the three selection task classes we assessed for each force-feedback condition. For right-handed interaction, Table 8 shows that participants when presented with no feedback move more quickly than with soft or hard force-feedback for all classes of selection tasks. Overall, a hard force-feedback condition produced poorest results, whereas soft force-feedback was better. For two-handed interaction, similar trends were less obvious to see.

5.3 Behaviour Graphs

As suggested from the time, distance and velocity results, participants use different paths to select multiple objects when under the three force-feedback conditions we assessed. To show this, we plotted a series of trajectory graphs for each condition. In Figure 6, we can see an example where there is a noticeable difference in how participants behaved. In this figure the coloured lines represent the trajectories from different feedback conditions: 'Blue Line' - hard force-feedback condition, 'Green Line' - soft force-feedback condition, 'Red Line' - no force-feedback condition. The coloured circles are the target spheres and they are shown at correct scale. In particular, for the no force-feedback condition, the inserts Ref 1 and Ref

Right-Handed Interaction			
Distance (m)			
	Select 1	Select 2	Select 3
Hard FB	0.18	0.26	0.39
Soft FB	0.17	0.27	0.38
No FB	0.20	0.44	0.66
Two-Handed Interaction			
Distance (m)			
	Select 1	Select 2	Select 3
Hard FB	0.21	0.38	0.48
Soft FB	0.18	0.33	0.46
No FB	0.25	0.60	0.88

Table 5: Average distance travelled to complete task

Right-Handed Interaction			
Distance (m)			
	Select 1	Select 2	Select 3
Soft FB vs Hard FB	-0.01	0.01	-0.01
Soft FB vs No FB	-0.03	-0.18	-0.27
Hard FB vs No FB	-0.01	-0.18	-0.26
Two-Handed Interaction			
Distance (m)			
	Select 1	Select 2	Select 3
Soft FB vs Hard FB	-0.03	-0.06	-0.02
Soft FB vs No FB	-0.07	-0.27	-0.42
Hard FB vs No FB	-0.04	-0.22	-0.39

Table 6: Average difference in distance travelled between each force-feedback condition tested

2 in Figure 6 highlight parts of trajectories that indicate that more time and distance was spent on the surface of a target as if to confirm its selection. In hard and soft force-feedback conditions this is not the case. To confirm this trend, we computed the time spent on the surface of each target for tasks involving 2 or more objects to select. As we can see from Table 9, on average participants when presented with a no force-feedback condition, spent the longest time in contact on the surface of targets before moving to the next object. Hard force-feedback responses produced the quickest times, whereas the results for soft force-feedback were more similar to the no feedback case, but again quicker.

In addition to this, we can see that the trajectories entering and leaving the surface of a target are again different for each force-feedback condition assessed. In cases with hard force-feedback, as shown in Figures 6 and 7, there are more 'arching' movements/deflections away from the target just after it is selected in comparison to when no force-feedback is present. In particular, we can see in the area labelled Ref 2 in Figure 6, the moment a target is selected with a hard force-feedback response, the next movement is away from the target. In contrast, with the soft force-feedback, this behaviour is less pronounced as participants can maintain their hand on the optimum path to the next target better.

5.4 Qualitative Data

From the qualitative data recorded as shown in Table 10, we can see participants found all force-feedback conditions ease to use and experienced little sickness. In terms of responsiveness and naturalism of interaction, conditions with haptic force-feedback produced best results, in particular hard force-feedback conditions. In contrast, participants found conditions with no feedback not natural or unresponsive to use and as a consequence were more aware of the outside surroundings, rather than the VE itself, indicating perhaps less presence was felt.

Right-Handed Interaction			
	Select 1	Select 2	Select 3
Hard FB vs Soft FB	4	0	2
No FB vs Hard FB	2	15	15
No FB vs Soft FB	7	15	15
Two-Handed Interaction			
	Select 1	Select 2	Select 3
Hard FB vs Soft FB	7	5	0
No FB vs Hard FB	11	15	15
No FB vs Soft FB	15	15	15

Table 7: Pair-wise comparison of distance travelled data: number of tasks out of 15 within selection class with significantly different results

Right-Handed Interaction			
Velocity (ms^{-2}) $\times 10^{-7}$			
	Select 1	Select 2	Select 3
Soft FB vs Hard FB	59.01	19.91	25.52
Soft FB vs No FB	-8.74	-8.28	-25.51
Hard FB vs No FB	-67.75	-28.18	-51.03
Two-Handed Interaction			
Velocity (ms^{-2}) $\times 10^{-7}$			
	Select 1	Select 2	Select 3
Soft FB vs Hard FB	1.00	6.54	20.51
Soft FB vs No FB	4.39	-1.40	27.11
Hard FB vs No FB	6.90	-7.94	6.60

Table 8: Average difference in velocity to complete task between each feedback condition tested

6 DISCUSSION

From the results, we have outlined instances where different types of haptic force-feedback can benefit the selection of 1, 2 and 3 targets. Using both the quantitative and qualitative data record, we can give the following summaries:

- No force-feedback: Allows participants to select with higher initial velocities. For simple selection tasks such as selecting only one target, results were similar to best times achieved under soft force-feedback conditions. However, as the difficulty of the task increases to the selection of 3 targets, participants performed the slowest and covered greater distances in comparison to soft or hard haptic force-feedback conditions.
- Soft force-feedback: Produced best results for time and distance taken to complete for all selection tasks assessed. In terms of velocities, participants performed at a similar level to when no force-feedback was present.
- Hard force-feedback: Produced performances similar to best distances achieved under soft force-feedback conditions. However, the velocity taken to cover these distances were the slowest out of the three force-feedback conditions assessed. As a result, the overall time taken was comparable to participants performing without any haptic force-feedback. However, when interacting with the right hand only, there were cases where participants performed very slowly.

To explain these variations, we argue that participants change how they move their hands depending on the type of force-feedback condition faced. By looking at the hand movement graphs again, we can see that under no force-feedback conditions, participants often travel more and spend a longer time on the surface of the intended target, 'dwelling' before moving on. Over multiple tasks, this behaviour manifests itself in slower completion times and larger distances covered by each hand. In comparison to hard force-feedback conditions, this behaviour does not exist, and participants, once in contact with the intended target move away quickly. Nevertheless,

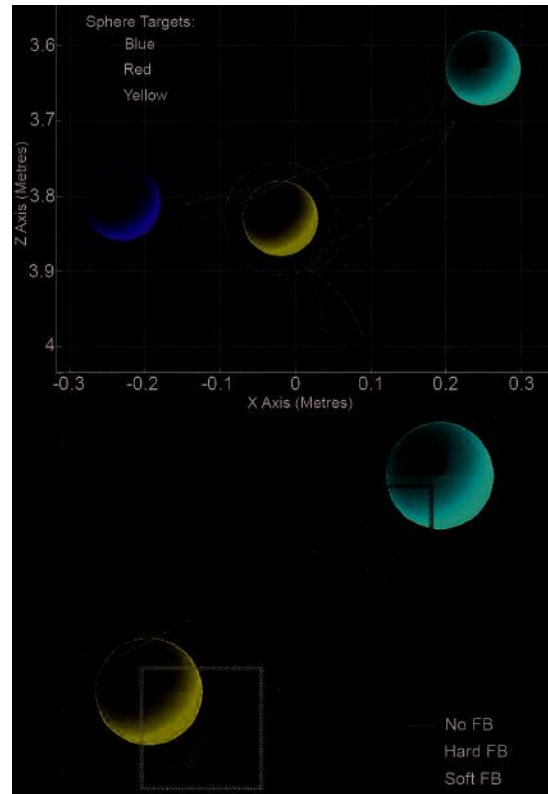


Figure 6: Example of trajectories for task number 34. See text for the meaning of the colours

Right-Handed Interaction			
	Time (s)		
	Select 2	Select 3	
	Blue to Red	Blue to Red	Red to Yellow
Hard FB	1.94	2.01	1.73
Soft FB	2.60	2.70	2.34
No FB	2.72	2.81	2.59

Table 9: Average time spent on surface of target before moving to another target for each feedback condition

even without this detrimental characteristic, the average time for task completion under hard force-feedback conditions was not better. To explain this, we can see from the trajectory graphs that participants often pull back and manoeuvre around objects and thus do not take the most optimum path between targets, effectively doing more work. Additionally the speed of movement is much slower. For the soft feedback case, this deflection and velocity characteristic is evident, though the timing and distances results are better.

7 CONCLUSION

In this paper, we explored how three different haptic force-feedback conditions affect user performance whilst selecting multiple targets. Overall, we found that soft force-feedback conditions produced best results in terms of time taken and distance travelled. For simple tasks, hard force-feedback responses led to results that were almost as good as soft force-feedback, but as the difficulty of the selection task increased, the differences in user performance between hard and no force-feedback conditions was less pronounced. By plotting the path trajectories for the three conditions we noticed that participants move their hands differently when presented with different types of haptic force-feedback. In particular, for no force-feedback

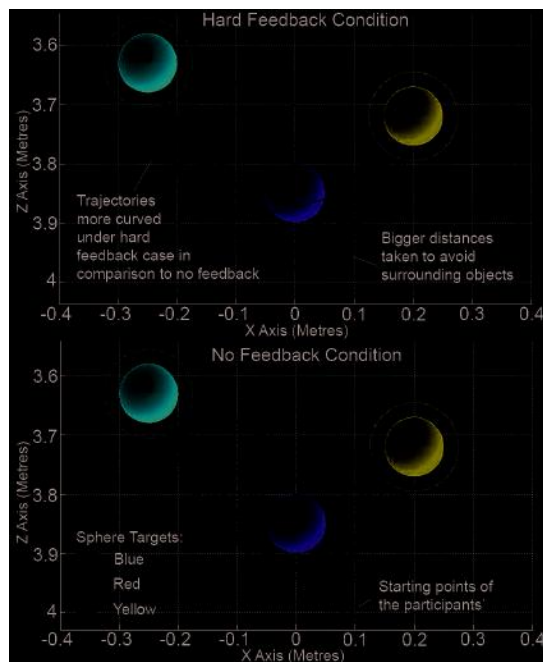


Figure 7: Example of trajectories for task number 38 (all subjects)

	Average Ranking (Out of 7)					
	Usability	Responsiveness	Naturalism	Awareness of Surroundings	Ease of Use	Sickness
No FB	5	4	3	4	5	1
Soft FB	6	5	4	3	5	2
Hard FB	6	7	7	2	6	2

Table 10: Average Ratings- usability and presence questionnaires

conditions, participants spend more time 'dwelling' on the surface of targets to confirm its selection ultimately adversely affecting their performance especially for multiple tasks. Even though this behaviour did not happen under hard feedback conditions, better user performances did not always follow. In fact, for conditions where there was hard force-feedback, we observed cases where participants moved more slowly and with greater 'arching' trajectories resulting in movements away from the optimum path to the next target. In general, this behaviour was minimised by reducing the hardness of the feedback, as seen less under soft feedback conditions. This contrasts with the user subjective ratings, where the hard feedback condition was preferred, perhaps because it gives more "tangible" feedback, and is more like touching a real object. We suggest that the trade-off of supporting physical properties versus providing fast task performance deserves more attention, especially in the situation of non-simple selection and manipulation tasks which is the direction of our future work.

REFERENCES

[1] C. Avizzano, S. Marcheschi, M. Angerilli, M. Fontana, M. Bergamasco, T. Gutierrez, and M. Mannegeis. A multi-finger haptic interface for visually impaired people. *Robot and Human Interactive Communication, 2003. Proceedings. ROMAN 2003. The 12th IEEE International Workshop on*, pages 165–170, Oct.-2 Nov. 2003.

[2] J. D. Boeck, E. Cuppens, T. D. Weyer, C. Raymaekers, and K. Coninx. Multisensory interaction metaphors with haptics and proprioception in virtual environments. In *NordiCHI '04: Proceedings of the third Nordic conference on Human-computer interaction*, pages 189–197, New York, NY, USA, 2004. ACM.

[3] J. D. Boeck, T. D. Weyer, C. Raymaekers, and K. Coninx. Using the non-dominant hand for selection in 3d. In *3DUI '06: Proceedings of*

the 3D User Interfaces, pages 53–58, Washington, DC, USA, 2006. IEEE Computer Society.

[4] D. A. Bowman, E. Kruijff, J. J. L. Jr., and I. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley, 2005.

[5] S. Choi and H. Z. Tan. Perceived instability of virtual haptic texture. ii. effect of collision-detection algorithm. *Presence: Teleoper. Virtual Environ.*, 14(4):463–481, 2005.

[6] C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon, and J. C. Hart. The cave - audio visual experience automatic virtual environment. *Commun. ACM*, 35(6):64–72, 1992.

[7] P. M. Fitts. The information capacity of the human motor system in controlling the amplitude of movement. pages 381–391, June 1954.

[8] A. Guerraz, C. Loscos, and H. Ritter Widenfeld. How to use physical parameters coming from the haptic device itself to enhance the evaluation of haptic benefits in user interface? 2006.

[9] A. E. Kirkpatrick and S. A. Douglas. Application-based evaluation of haptic interfaces. In *HAPTICS '02: Proceedings of the 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, page 32, Washington, DC, USA, 2002. IEEE Computer Society.

[10] D. A. Lawrence, L. Y. Pao, A. M. Dougherty, M. A. Salada, and Y. Pavlou. Rate-hardness: a new performance metric for haptic interfaces. *Robotics and Automation, IEEE Transactions on*, 16(4):357–371, 2000.

[11] S. Lee and G. J. Kim. Effects of haptic feedback, stereoscopy, and image resolution on performance and presence in remote navigation. *Int. J. Hum.-Comput. Stud.*, 66(10):701–717, 2008.

[12] D. B. Manek. Effects of visual displays on 3d interaction in virtual environments. Blacksburg, Va.: University Libraries, Virginia Polytechnic Institute and State University, 2004.

[13] I. Oakley, S. Brewster, and P. Gray. Solving multi-target haptic problems in menu interaction. In *CHI '01: CHI '01 extended abstracts on Human factors in computing systems*, pages 357–358, New York, NY, USA, 2001. ACM.

[14] I. Oakley, M. R. McGee, S. Brewster, and P. Gray. Putting the feel in 'look and feel'. In *CHI '00: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 415–422, New York, NY, USA, 2000. ACM.

[15] W. R. Provancher, M. R. Cutkosky, K. J. Kuchenbecker, and G. Niemeyer. Contact location display for haptic perception of curvature and object motion. *Int. J. Rob. Res.*, 24(9):691–702, 2005.

[16] K. Rasmus-Gröhn, C. Magnusson, and H. E. Efring. Iterative design of an audio-haptic drawing application. In *CHI '07: CHI '07 extended abstracts on Human factors in computing systems*, pages 2627–2632, New York, NY, USA, 2007. ACM.

[17] E. Ruffaldi, A. Frisoli, M. Bergamasco, C. Gottlieb, and F. Tecchia. A haptic toolkit for the development of immersive and web-enabled games. In *VRST '06: Proceedings of the ACM symposium on Virtual reality software and technology*, pages 320–323, New York, NY, USA, 2006. ACM.

[18] E. Samur, F. Wang, U. Spaelter, and H. Bleuler. Generic and Systematic Evaluation of Haptic Interfaces Based on Testbeds. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'07)*, 2007.

[19] A. Steed. Towards a general model for selection in virtual environments. In *3DUI '06: Proceedings of the 3D User Interfaces*, pages 103–110, Washington, DC, USA, 2006. IEEE Computer Society.

[20] D. Swapp, V. Pawar, and C. Loscos. Interaction with co-located haptic feedback in virtual reality. *Virtual Reality*, 10(1):24–30, 2006.

[21] L. Vanacken, C. Raymaekers, and K. Coninx. Evaluating the influence of multimodal feedback on egocentric selection metaphors in virtual environments. In *HAID*, pages 12–23, 2006.

[22] S. Wall, K. Paynter, A. Shillito, M. Wright, and S. Scali. The effect of haptic feedback and stereo graphics in a 3d target acquisition task. In *Proceedings of Eurohaptics 2002, University of Edinburgh, 8-10th July*, pages 23–29. N/A, 2002.

[23] S. A. Wall and W. S. Harwin. Quantification of the effects of haptic feedback during a motor skills task in a simulated environment. In *In Proceedings at Phantom User Research Symposium00*, pages 61–69, 2000.